



## *Triaging a Database Slowdown*

NYOUG

December 6, 2007

*Michael S. Abbey — Database Officer Oracle Practice*



## *What You Need to Get Started*

- Patience
- Diplomacy
- Warmth
- Kindness
- Tolerance
- Sense of humour
- Same old same old
- It's probably the CxO
- Swaddle your users
- Positive outlook
- High threshold
- Lighten up the dialogue



## *No no no no no ...*

- Why are you calling me again; it was not the database the last 2,328,281,000,029,182,727 times!
- If I can find the time, I will check it out ...
- Did you check the app-tier?
- Don't look at me; it's those b\_\_\_\_y developers.
- Call Miguel or Honorée

3

*Michael S. Abbey — Triaging a Database Slowdown*



## *What You Really Need*

- Metrics (if possible)
  - is the load average "high"
  - is the hourglass there a lot longer than usual?
  - on a scale of 1 to 10, how bad is the slowdown
- What is the caller doing with what system when the slowdown is noticed?


4

*Michael S. Abbey — Triaging a Database Slowdown*



## *Help Translate*

- 'There's something wrong with the database



- Can you work with development, if necessary, to find out why the applications are so slow?

5

*Michael S. Abbey — Triaging a Database Slowdown*



## *Scenario*

- CPU maxed out
- Queries seem to take forever
- Reporting performance is noticeably slower than usual
- Taking forever to log in
- Hourglass city



6

*Michael S. Abbey — Triaging a Database Slowdown*



# UNIX

vmstat 2 15 : system/user/idle CPU

```
oracle@ntirety-dbia-->(tket) ** Production **
/home/oracle> vmstat 2 15

procs-----memory-----swap-----io-----system-----cpu-----
r  b   swpd   free   buff   cache   si   so    bi   bo    in   cs   us   sy   id   wa
2  0     0 38192 431356 2475440 0    0   602  109    2    1  17  7  72  4
2  0     0 37424 432136 2475408 0    0   382  130 1099   733  44  6  49  0
4  0     0 34192 432728 2475428 0    0   290  212 1099 19664  62 19 19  0
3  0     0 29424 433488 2475484 0    0   354  204 1173 27854  51 25 24  0
14 0     0 20072 434236 2475416 0    0   322  806 1278 14252  47 16 37  0
10 0     0 28512 345840 2473712 0    0  420 1952 1464  2567  72 28  0  0
5  0     0 63056 346320 2473708 0    0  246  410 1120  649  87 15  0  0
6  0     0 78608 346660 2473776 0    0   182  384 1111  615  88 12  0  0
4  0     0 98920 347180 2473800 0    0  284  356 1129  573  92  8  0  0
8  2     0 98320 347836 2473892 0    0  258  882 1342  7330  90 10  0  0
9  0     0 95440 348496 2473980 0    0  192 1856 1556 15816  83 16  0  0
9  0     0 94808 348976 2474044 0    0  246  324 1127 12810  86 13  0  0
7  0     0 92512 349604 2474028 0    0  290  452 1167 12445  86 14  0  0
3  1     0 91072 350224 2474632 0    0  512  782 1293  4790  92  6  0  1
5  0     0 92568 350956 2486820 0    0  6346 1478 1501 1591  93  6  0  1
4  0     0 103120 351488 2486900 0    0  260  282 1113  794  90  6  4  0
3  0     0 102936 352204 2486932 0    0  352  156 1107 1225  74  6 20  0
3  0     0 102104 353032 2486920 0    0  418  104 1098 1425  73  2 25  0

oracle@ntirety-dbia-->(tket) ** Production **
/home/oracle>
```

7

Michael S. Abbey — Triaging a Database Slowdown



top (topas AIX): who are the biggest consumers

```
VPN-dbia - SecureCRT
top - 09:44:20 up 64 days, 20:59, 3 users, load average: 4.46, 4.70, 4.86
Tasks: 237 total, 3 running, 234 sleeping, 0 stopped, 0 zombie
Cpu(s): 46.8% us, 9.0% sy, 0.0% ni, 40.9% id, 0.2% wa, 0.0% hi, 3.1% si
Mem: 4040456k total, 3973832k used, 66624k free, 410684k buffers
Swap: 0k total, 0k used, 0k free, 2452524k cached


  PID USER   PR   NI  VIRT  RES  SHR  S  %CPU  %MEM    TIME+  COMMAND
  946  root    17    0  3780   652  260  R  99.9   0.0 384:51.33  gzip
13573 oracle  25    0  941m 297m 288m  R  99.9   7.5  89:51.15  oracle
2864  oracle  16    0 109m  9.9m 6232  S 17.6   0.3  8224:19  tnslnsr
3353  oracle  16    0  571m  46m 4208  S 14.6   1.2 10132:54  opmn
20825 oracle  16    0  962m  39m 15m   S  2.0   1.0  0:37.65  oracle
18859 oracle  16    0  7304 1180  784  R  0.3   0.0  0:01.47  top
  1  root    16    0  5776  588  484  S  0.0   0.0  1:58.65  init
  2  root    RT    0  0  0  0  S  0.0   0.0  0:18.47  migration/0
  3  root    34  19  0  0  0  S  0.0   0.0 10:25.66  ksoftirqd/0
  4  root    RT    0  0  0  0  S  0.0   0.0  0:19.33  migration/1
  5  root    34  19  0  0  0  S  0.0   0.0 16:14.22  ksoftirqd/1
  6  root    RT    0  0  0  0  S  0.0   0.0  0:16.87  migration/2
  7  root    34  19  0  0  0  S  0.0   0.0  7:42.29  ksoftirqd/2
  8  root    RT    0  0  0  0  S  0.0   0.0  0:16.85  migration/3
  9  root    34  19  0  0  0  S  0.0   0.0  9:21.78  ksoftirqd/3
 10  root    5 -10  0  0  0  S  0.0   0.0  0:01.15  events/0
 11  root    5 -10  0  0  0  S  0.0   0.0 55:33.11  events/1
 12  root    5 -10  0  0  0  S  0.0   0.0  0:00.84  events/2
```

8

Michael S. Abbey — Triaging a Database Slowdown



## Windows



9

*Michael S. Abbey — Triaging a Database Slowdown*



# A Locking Problem

10

*Michael S. Abbey — Triaging a Database Slowdown*



## *In General 10g/11g*

```
select sid, blocking_session,  
       username,  
       blocking_session_status  
from v$session  
order by blocking_session;
```

[Best MetaLink document](#)

11

*Michael S. Abbey — Triaging a Database Slowdown*



## *In General*

```
SQL> select * from dba_blockers;  
HOLDING_SESSION
```

-----  
**1201**

```
SQL> desc dba_blockers  
Name                               Type
```

```
-----  
HOLDING_SESSION    NUMBER
```

12

*Michael S. Abbey — Triaging a Database Slowdown*



## *In General*

```
SQL> desc dba_waiters
```

| Name            | Type         |
|-----------------|--------------|
| -----           | -----        |
| WAITING_SESSION | NUMBER       |
| HOLDING_SESSION | NUMBER       |
| LOCK_TYPE       | VARCHAR2(26) |
| MODE_HELD       | VARCHAR2(40) |
| MODE_REQUESTED  | VARCHAR2(40) |
| LOCK_ID1        | NUMBER       |
| LOCK_ID2        | NUMBER       |

1201

13

*Michael S. Abbey — Triaging a Database Slowdown*



## *So, What is 1201 Up To?*

```
select sql_text
  from v$sqlarea
 where address =
        (select sql_address
          from v$session
         where sid = 1201);
```

14

*Michael S. Abbey — Triaging a Database Slowdown*



## *Caveats About That Query*

- There may be more than 1 SQL statement associated with SID 1201
  - *single row subquery returns more than one row*
  - may have to change "=" to "in"
- May return *no rows selected*



15

*Michael S. Abbey — Triaging a Database Slowdown*



## *Armed With / Can Get*

- SID of blocking session
- SID of waiting session(s)
- SQL text associated with that blocking session
- SQL text associated with waiting session(s)
- Information about the environment of the blocking session from V\$SESSION
  - USERNAME MACHINE OSUSER

16

*Michael S. Abbey — Triaging a Database Slowdown*





# Bad SQL

(more long term)

17

*Michael S. Abbey — Triaging a Database Slowdown*



## *The Developers' Treaty (1983)*

The undersigned agree that neither party will assess any blame on one another, make any negative comments about each other's code, nor make any remarks that are intended to or may be construed by a third party to be critical of said parties' code; furthermore ...

Section 3-D; subsection 12; exhibit 34-C12

18

*Michael S. Abbey — Triaging a Database Slowdown*



## *Bad SQL*

- Inefficient logical reads
  - can be more expensive than physical reads
  - physical reads deemed "too expensive" before advances in storage technology
- Logical vs. physical
  - high buffer gets equates to the hourglass
  - minimize (almost at all costs)

19

*Michael S. Abbey — Triaging a Database Slowdown*



## *The Plague of the Oracle Database*

**Buffer busy waits** occur when a session wants to access a database block in the buffer cache but it cannot as the buffer is busy. The two main cases where this can occur are:

- another session is reading the block into the buffer
- another session holds the buffer in an incompatible mode to a new request

20

*Michael S. Abbey — Triaging a Database Slowdown*



## *The Plague of the Oracle Database*

**Latch free waits** occur when a process waits for a latch that is currently busy (held by another process); the wait time increases exponentially.

- in willing-to-wait mode a process will spin and try to get the same latch again
- in no-wait mode a process will "give-up" and request a different latch if one is unavailable

21

*Michael S. Abbey — Triaging a Database Slowdown*



## *Run 12 Times in 2 Minutes*

```
select event,count(*)
  from v$session_wait
 where event in ('latch free',
                'buffer busy waits')
 group by event;
```

22

*Michael S. Abbey — Triaging a Database Slowdown*



## *Just Remember*

- These 2 wait events are the nature of the beast with Oracle
- They appear at least 50% (if not more) of the time in *v\$sqlsession\_wait*
- Persistent double-digits for **one** or **both** over an interactive running of the query may be cause for concern

23

*Michael S. Abbey — Triaging a Database Slowdown*



```
latch free 12  buffer busy waits  2
latch free 14
latch free  6
latch free  9  buffer busy waits  1
latch free  2  buffer busy waits  1

latch free  1  buffer busy waits  9
latch free  9  buffer busy waits  2

latch free  2  buffer busy waits  2
latch free  2
```

24

*Michael S. Abbey — Triaging a Database Slowdown*



```
latch free 2    buffer busy waits 2
latch free 1

                buffer busy waits 1
                buffer busy waits 1

latch free 1

                buffer busy waits 2
latch free 2
```

25

*Michael S. Abbey — Triaging a Database Slowdown*



## *Longer Term Triage*

- Where and what is the hot data
  - querying the data dictionary
  - turning on table monitoring (SYS.DBA\_TAB\_MODIFICATIONS)
- How's my buffer cache
  - bigger is not always better
  - should I have multiple block-sized buffer cache



26

*Michael S. Abbey — Triaging a Database Slowdown*



## *Where?*

```
select d.tablespace_name,  
       sum(x.count) total_waits,  
       sum(x.time) time_waited  
from sys.x$kcbbwait x,  
     sys.dba_data_files d  
where x.inst_id = userenv('Instance')  
      and x.count > 0  
      and x.indx + 1 = d.file_id  
group by d.tablespace_name  
order by 3 desc;
```

27

*Michael S. Abbey — Triaging a Database Slowdown*

## *What's Happening?*

```
select decode(state, 0, 'Free',  
              1, 'Read and Modified',  
              2, 'Read and Not Modified',  
              3, 'Currently Being Read',  
              'Other' ), count(*)  
from sys.x$bh  
group by decode(state, 0, 'Free',  
                1, 'Read and Modified',  
                2, 'Read and Not Modified',  
                3, 'Currently Being Read',  
                'Other' )
```

28

*Michael S. Abbey — Triaging a Database Slowdown*



## *Finding Culprits*

- Buffer gets (STATSPACK)
- Consistent gets (AUTOTRACE)
- sqlplus
  - *explain plan*
  - *set autot trace exp*
- Today's small problems are tomorrow's disasters



29

*Michael S. Abbey — Triaging a Database Slowdown*



## *Lessons to be Learned*

- Your less-technical colleagues will be looking for a quick fix
- A "well-behaved" database all of a sudden performs poorly may be a contributor but not the culprit
- Queries may start behaving differently as data volumes change ... **hints and the cost-based optimizer**

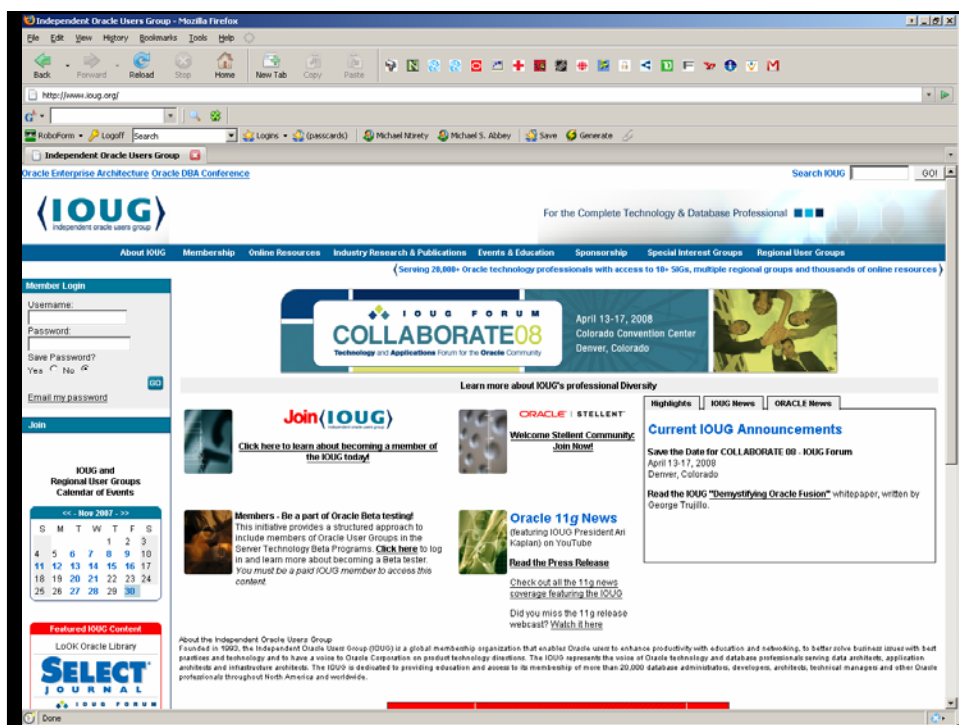
30

*Michael S. Abbey — Triaging a Database Slowdown*



31

Michael S. Abbey — Triaging a Database Slowdown







michael.abbey@ntirety.com



**YAHOO!**



fenderpbs

