

OraPub is about Oracle performance.



- OraPub is all about Oracle performance management; monitoring, firefighting, quantitative and predictive analysis.
- Web site started in 1995 and the company was founded in 1998 by Craig Shallahamer.
- OraPub has always been about disseminating Oracle database centric technical information.
- Consulting, training, books, papers, and products are now being offered.
- We have been on-site in 23 countries and our resources have been received in probably every country where there are DBAs.

Resources

- Free Tools
- Free Papers
- Books
- Products
- Consulting
- Training


Short resume.

- Studies economics, mathematics, and computer science at university in California, US.
- Started working with Oracle technology in 1989 as a Forms 2.3 developer on Oracle version 5.
- Soon after started performance firefighting...daily!
- Co-found both Oracle's Core Technology and System Performance Groups.
- Left Oracle to start OraPub, Inc. in 1998.
- Authored 24 technical papers and worked in 24 countries.
- Authors and teaches his classes *Oracle Performance Firefighting*, *Adv Oracle Performance Analysis*, and *Oracle Forecasting & Predictive Analysis*.
- Authored the books, *Forecasting Oracle Performance*, and his new book, *Oracle Performance Firefighting*.



If you want my firefighting book...

- Retail \$49.95
- Amazon \$69.95
- OraPub normal \$39.95
- "FLORIDA" disc. \$29.96
- Cash in hand \$30.00



©2010 OraPub, Inc.

A day in the life of an Oracle server process



Hurry up and wait .

```
SQL> 1
  1 select event doings,
  2         time_waited/100 time_s
  3 from   v$session_event
  4 where  sid = &sid
  5       and time_waited > 0
  6 union
  7 select 'burning CPU' doings,
  8        value/1000000 time_s
  9 from   v$sess_time_model
 10 where  sid=&sid
11*    and stat_name = 'DB CPU'
SQL> /
```

Doings	TIME_S
SQL*Net message from client	633.550
burning CPU	0.088
db file scattered read	0.010

3 rows selected.



Being probed is not pleasant!

```
[oracle@fourcore SEOUC]$ ps -eaf|grep 3635
oracle    3635   3634   0 16:35 ?        00:00:00 oracleprod18
oracle    9314   9268   0 16:51 pts/1    00:00:00 grep 3635

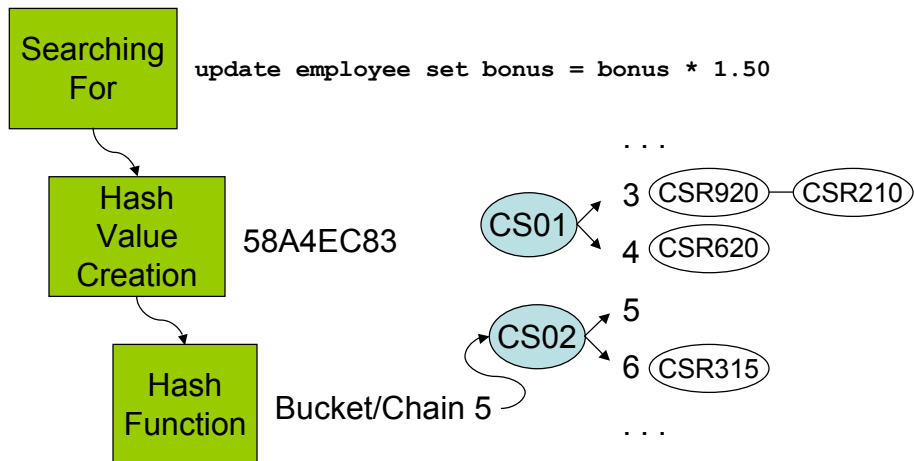
[oracle@fourcore SEOUC]$ strace -p 3635
Process 3635 attached - interrupt to quit
read(9,
```



©2010 OraPub, Inc.


**Client expectations are
so unbelievably high!**

I'm very environmentally conscious.




©2010 OraPub, Inc.

If only life was so simple...



©2010 OraPub, Inc.

But it's not that simple!



©2010 OraPub, Inc.

I've got to "ask permission."

```
Function Get_Mutex(mutex_name,type,mode)
{
  If type eq 'nowait'
    If Fast_Get(mutex_name)
      return TRUE
    Else
      return FALSE
    End-If
  Else
    If Fast_Get(mutex_name)
      Then
        return TRUE
      Else
        While ( TRUE )
          If Spin_Get(mutex_name)
            Then
              return TRUE
            Else
              Register_Event("cursor: *")
              Mutex_Wait(try++)
            End-If
          End-While
        End-If
      End-If
    }
  }

Function Mutex_Wait(try)
{
  options defined by kernel developers
  - CPU yield, blocking wait, sleep
}
```

```
Function Fast_Get(mutex_name)
{
  If mutex.holder := sid occurs
    Then
      Case mode:
        'X': If mutex.ref_count = 0
          Then
            return TRUE
          Else
            mutex.holder = clear
            return FALSE
          End-If
        'S': mutex.ref_count++
            mutex.holder = clear
            return TRUE
      End-Case
    Else
      return FALSE
    End-If
  }

Function Spin_Get(mutex_name)
{
  for i = 1 to 255
    If Fast_Get(mutex_name)
      Then
        return TRUE
      End-If
    End-For
  return FALSE
}
```

©2010 OraPub, Inc.

Burn and wait

```
SQL> 1
      2 select event doings,
      3         time_waited/100 time_s
      4 from   v$session_event
      5 where  sid = &sid and time_waited > 0
      6 union
      7 select 'burning CPU' doings,
      8         value/1000000 time_s
      9 from   v$sess_time_model
      0 where sid=&sid and stat_name = 'DB CPU'
      1 order by time_s desc
SQL> /
```

Doings	TIME_S
SQL*Net message from client	763.252
burning CPU	130.955
library cache: mutex X	7.230
latch: shared pool	1.660
db file sequential read	0.070
db file scattered read	0.010

6 rows selected.



©2010 OraPub, Inc.

Cursor is not found! ...hard parse

First, I've got to get some memory.


SP #1



SP #2



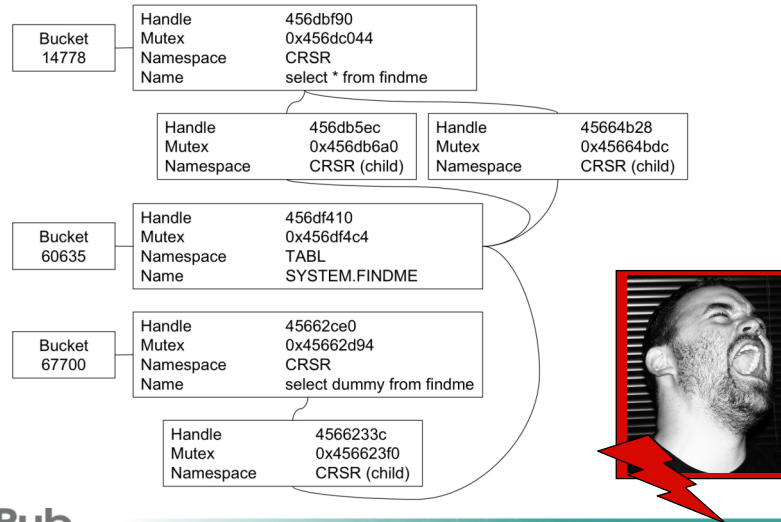
SP #3



©2010 OraPub, Inc.


Cursor is not found! ...hard parse

Second, I've got to establish establish relationships.




©2010 OraPub, Inc.

The cascade effect is real.



©2010 OraPub, Inc.

Especially when sharing occurs.



Time to run the SQL!

First, pin the cursor.

Second, find out where the block is.

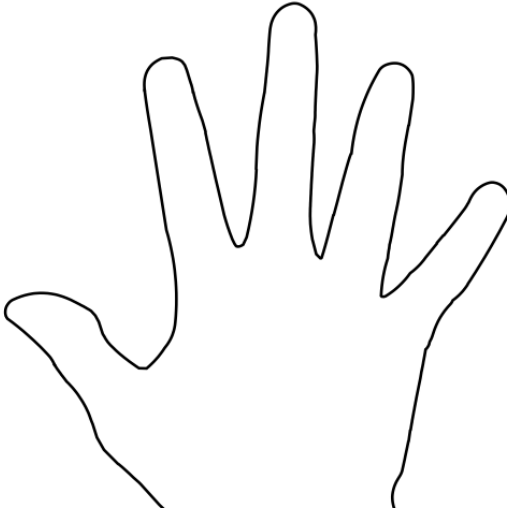
Third, get the block and pin it.

Fourth, change the row/block.

Fifth, generate some redo, undo, redo.

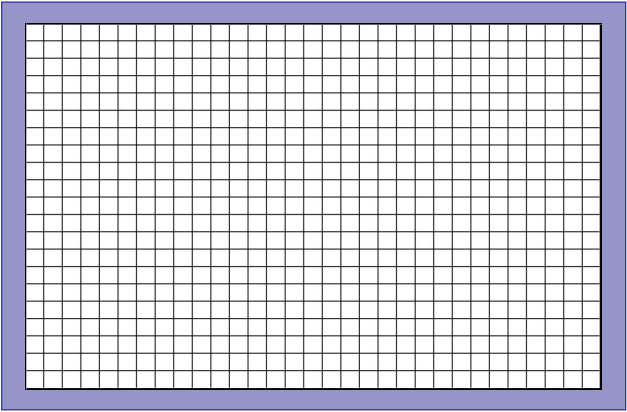
Sixth, let my client know I'm done.

Pin the cursor.




©2010 OraPub, Inc.

Find out where the block is.



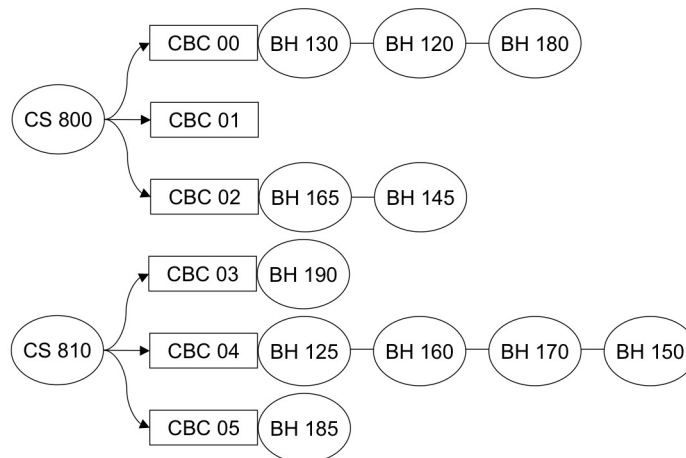
©2010 OraPub, Inc.

Buffer header close-up

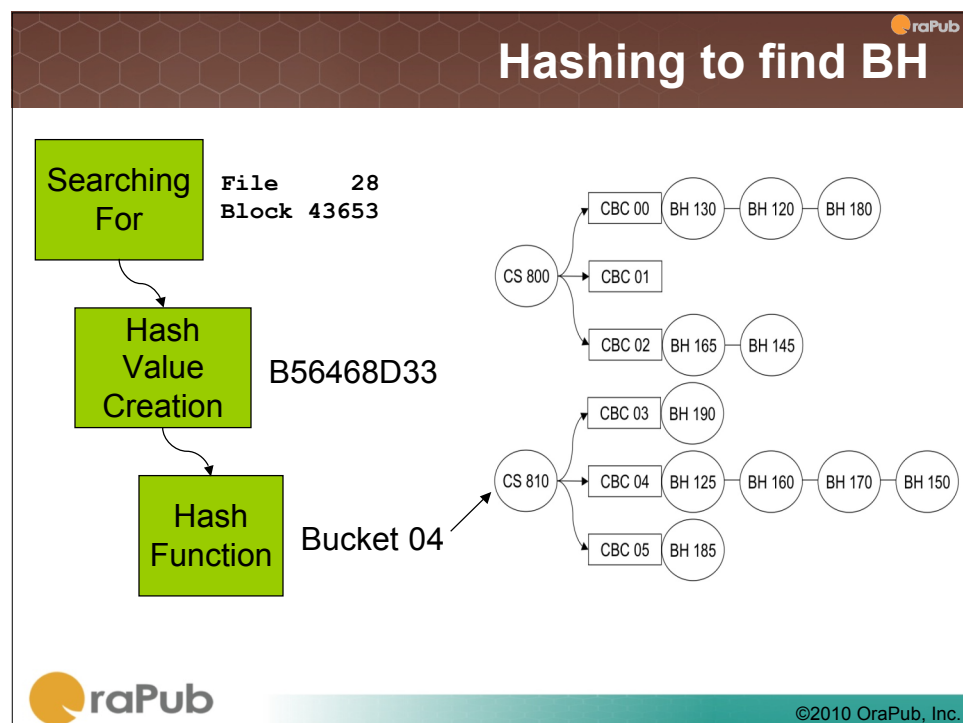



©2010 OraPub, Inc.

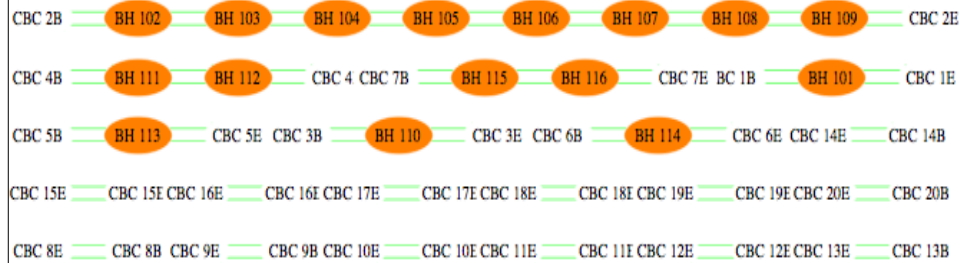
A man's got to have a good map!



©2010 OraPub, Inc.



Popular or long chains are problematic.



©2010 OraPub, Inc.

CBC latch contention

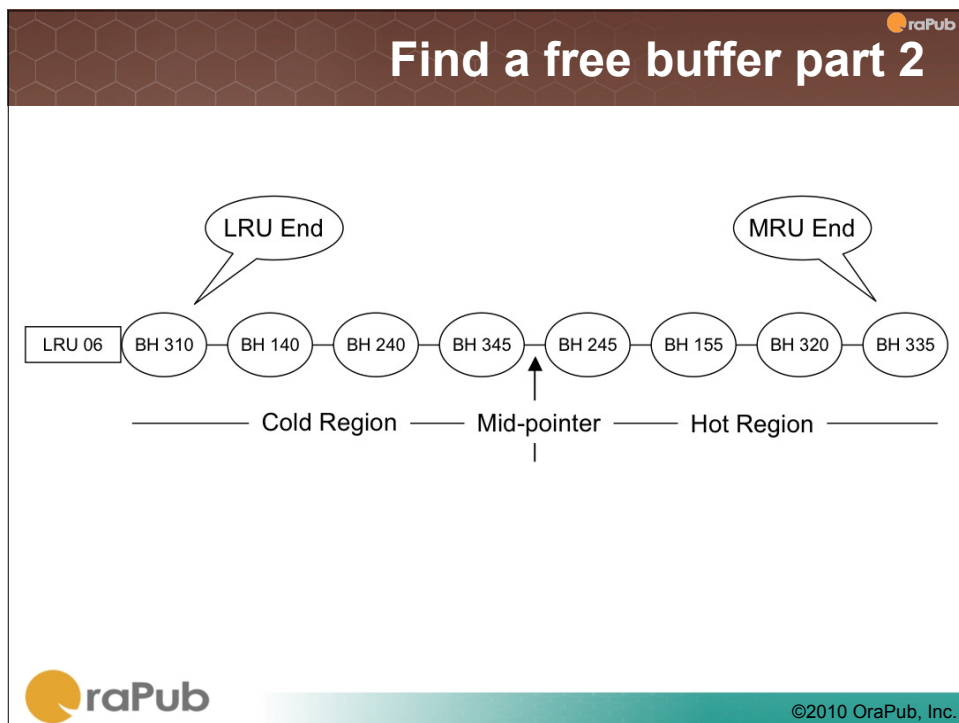

```
SQL> 1
2   select event doings,
3         time_waited/100 time_s
4   from   v$session_event
5  where   sid = &sid and time_waited > 0
6  union
7  select 'burning CPU' doings,
8        value/1000000 time_s
9  from    v$sess_time_model
10 where   sid=&sid and stat_name = 'DB CPU'
11* order by time_s desc
SQL> /
```

Doings	TIME_S
SQL*Net message from client	962.228
burning CPU	341.065
latch: cache buffers chains	9.230
db file scattered read	0.570
db file sequential read	0.321
read by other session	0.312

6 rows selected.



©2010 OraPub, Inc.



Get the block from “disk”



```
[oracle@fourcore ~]$ strace -rp 15621
Process 15621 attached - interrupt to quit
. . .
0.000041 gettimeofday({1266354967, 147052}, NULL) = 0
0.000044 pread64(10, "\6\242\0\0a. . .1\0\3\246~\1"... , 8192, 766255104) = 8192
0.022104 gettimeofday({1266354967, 147200}, NULL) = 0
. . .
```



©2010 OraPub, Inc.

I'm not a selfish person!



```
SQL> l
 1 select event doings,
 2         time_waited/100 time_s
 3 from   v$session_event
 4 where  sid = &sid and time_waited > 0
 5 union
 6 select 'burning CPU' doings,
 7        value/1000000 time_s
 8 from   v$sess_time_model
 9 where  sid=&sid and stat_name = 'DB CPU'
10* order by time_s desc
SQL> /
```

Doings	TIME S
SQL*Net message from client	1231.286
burning CPU	413.591
direct path read	392.073
read by other session	69.612
db file sequential read	8.721
latch: cache buffers chains	7.230

6 rows selected.



©2010 OraPub, Inc.

Update the row

```
SQL> alter system dump datafile 1 block 75847;

System altered.
SQL> !cat prod5_ora_21741.trc
...
Block header dump: 0x00412847
Object id on Block? Y
seg/obj: 0xff6b csc: 0x00.50fcb6 itc: 3 flg: 0 typ: 1 - DATA
fsl: 0 fnx: 0x412848 ver: 0x01

  Itl          Xid          Uba          Flag    Lck         Scn/Fsc
0x01    0x0003.00d.00000318  0x00c3e3d0.0593.0c  ----     1    fsc 0x0000.00000000
0x02    0x0008.01b.00000340  0x00c41bce.0481.24  ----     1    fsc 0x0000.00000000
0x03    0x0001.000.00000320  0x00c45fa0.0599.0b  ----     1    fsc 0x0000.00000000
... .
```



©2010 OraPub, Inc.

Let's not forget about redo, undo, redo.

```
SQL> l
 1  select event doings,
 2         time_waited/100 time_s
 3  from   v$session_event
 4  where  sid = &sid and time_waited > 0
 5  union
 6  select 'burning CPU' doings,
 7         value/1000000 time_s
 8  from   v$sess_time_model
 9  where  sid=&sid and stat_name = 'DB CPU'
10* order by time_s desc
SQL> /

Doings                                TIME_S
-----
SQL*Net message from client           1231.286
burning CPU                           490.890
log buffer space                      196.219
log file parallel write*              71.483
db file parallel write*               42.309
direct path write                     19.756

6 rows selected.
```



©2010 OraPub, Inc.

Let my client know I'm done!

```
SQL> 1
  2  select event doings,
  3         time_waited/100 time_s
  4  from   v$session_event
  5  where  sid = &sid and time_waited > 0
  6  union
  7  select 'burning CPU' doings,
  8         value/1000000 time_s
  9  from   v$sess_time_model
 10  where  sid=&sid and stat_name = 'DB CPU'
10* order by time_s desc
SQL> /
```

Doings	TIME_S
SQL*Net message from client	1432.589
burning CPU	499.691
log buffer space	196.219
log file parallel write	71.483
db file parallel write	42.309
direct path write	19.756
SQL*Net message to client	0.086

7 rows selected.

And life is boring once again!

```
SQL> 1
  2  select event doings,
  3         time_waited/100 time_s
  4  from   v$session_event
  5  where  sid = &sid and time_waited > 0
  6  union
  7  select 'burning CPU' doings,
  8         value/1000000 time_s
  9  from   v$sess_time_model
 10  where  sid=&sid and stat_name = 'DB CPU'
10* order by time_s desc
SQL> /
```

Doings	TIME_S
SQL*Net message from client	2576.002
burning CPU	521.621
log buffer space	196.219
log file parallel write	71.483
db file parallel write	42.309
direct path write	19.756
SQL*Net message to client	0.086

7 rows selected.

A day in the life of an Oracle server process



Want to dig deeper?

- **Training** from OraPub
 - Oracle Performance Firefighting (I)
 - Adv Oracle Performance Analysis (II)
 - Oracle Forecasting & Predictive Analysis
 - OraPub 1-Day 2010 Perf Seminar
- **Books**
 - Oracle Performance Firefighting (C. Shallahamer)
 - Forecasting Oracle Performance (C. Shallahamer)
- **Products**
 - **OraPub Stress Identifier** now available.
- **Papers/Tools** at www.orapub.com
 - Scientifically Evaluating Alternative Performance Solutions
 - Introduction To Oracle Performance Optimization
 - **Exploring Oracle's [Buffer Cache, Library Cache]**
- **Craig's Blog** at blogspot.com

**Irvine, CA
November**

**Philly, PA
October!!**

Tomorrow!!



©2010 OraPub, Inc.

A day in the life of an Oracle server process

