What if Kramer was your DBA & Seinfeld Tuned your Database?

New York - 2011

"Jerry, I can't find my backup"

Rich Niemiec, Rolta TUSC (www.tusc.com)`

Audience Knowledge

- Oracle9*i* Experience ?
- Oracle9*i* RAC Experience?
- Oracle10*g* Experience?
- Oracle Database 11g Experience

- Goals
  - Tuning Tips & AWR / Grid Tuning
  - Focus on a few nice features of Oracle 10g & 11g
- Non-Goals
  - Learn ALL aspects of Tuning Oracle

Overview

- Kramer's missing Backup
- George's Untuned System
- Elaine's Untouched System
- Jerry's Perfect Tuning Plan
- Statspack / AWR
  - Top Waits
  - Load Profile
  - Latch Waits
  - Top SQL
  - Instance Activity
  - File I/O
- The Future EM & ADDM
- Helpful V$/X$
- Summary

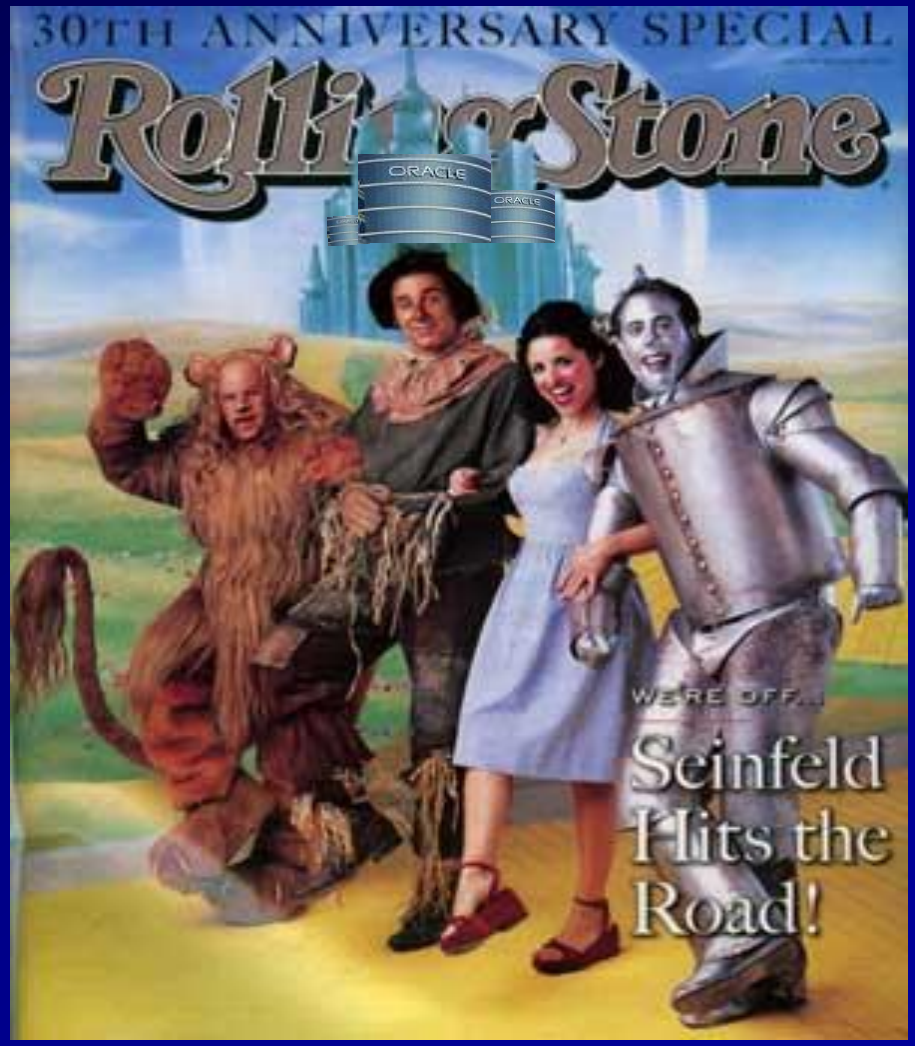

Kramer doesn't have a Backup

My Junior DBA is getting the backup right now!

What Kramer did…

- He was logged into production vs. test
- He deleted some production data
- His backup tape was at Jerry's apartment
- He taped a Lady Gaga song over the backup tape.
- He never actually tested the backup so the older backup tapes don't work either
- He doesn't have a DR site

Jerry reminds Kramer what he could have done to prevent all of this...

- The backup should have been in a secure location
- With 10g or 11g **encrypt the backup** so it will always be protected
- Could have used Oracle's Flashback and get deleted data back
- Data Guard allows you to fail over to a new sight.
- **Test your recovery & DR** to ensure it will

Jerry reminds Kramer what he could have done to prevent all of this...

- Just because a database may need to be recovered, do not delete the "corrupted" database *if possible*.
  - First, take a backup taken of the "corrupt" database.
  - If the restore does not work and you did not backup the "corrupt" database, you may have nothing to work with.
- When trouble shooting a problem query never let operations reboot the instance.
- Never startup a standby database in normal mode.

How to avoid this mistake...

- Kramer could have ensured that he never made this mistake in the first place by setting up his prompt to show exactly where he was. For unix/linux, put in your .profile, or .bash_profile.

If using the bash shell:

```
export PS1="[$(whoami)@$(hostname)]$ "
```

For sqlplus, it is the sqlprompt setting:

```
MIKE6:system:mike61> show sqlp
sqlprompt "MIKE6:system:mike61> "
```

It is setup in a sqlplus admin file:

```
[oracle@mikek5 oracle]$ tail -2
    $ORACLE_HOME/sqlplus/admin/glogin.sql
```

George Doesn't Tune Anything

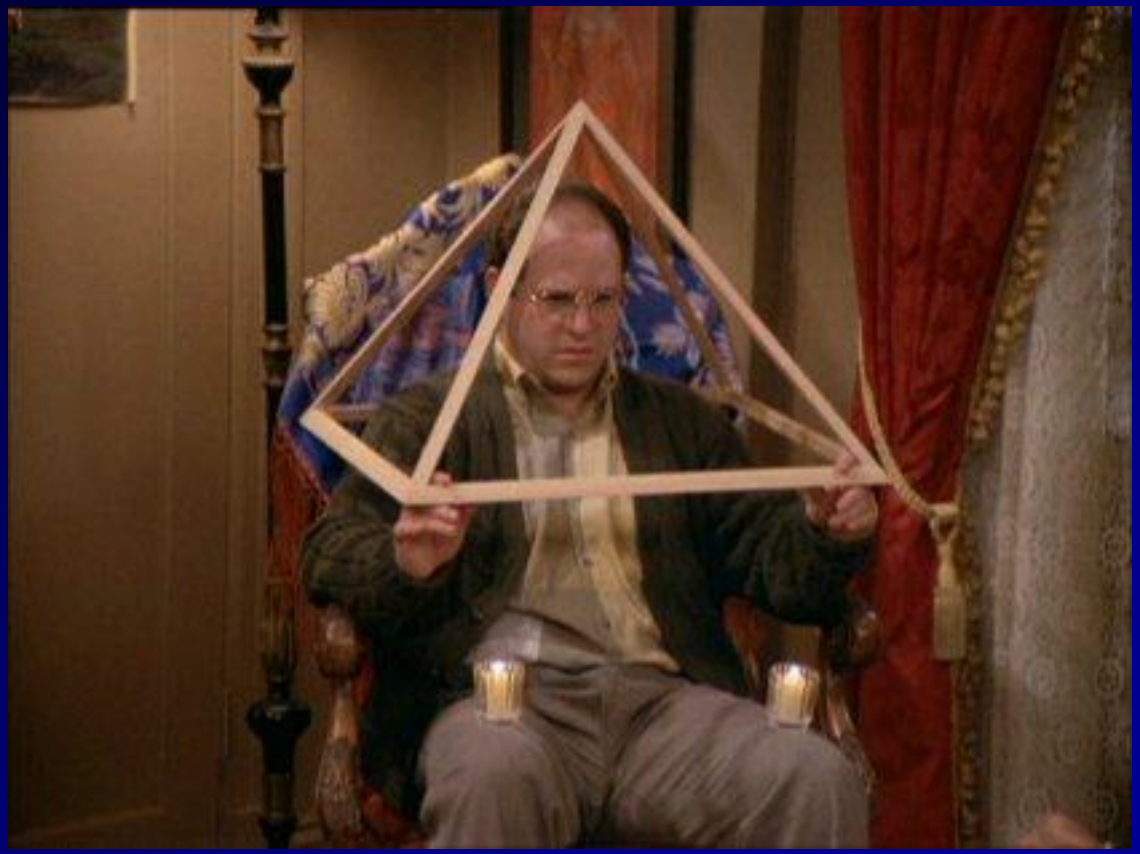

Instead of proactively tuning are you depending on other means to save you?

George doesn't Tune Anything...

- George doesn't believe in backups; It slows down the system.
- He uses the "kill -9" for anything slow
- George doesn't patch things especially security
- He uses default passwords for speed and so he doesn't have to change any application code.
- He tries not to do anything that requires actual work to be done
- He never tells anyone that he's going to bring down the system, he just does a "oops, it crashed" with a Shutdown Abort when he needs to bring it down

Jerry's Advice to George...

- **Default passwords should be changed** at database creation time.
- Good design beats ad-hoc design.
  - Don't work in a black box when tuning.
  - Establish priorities and work on what is important to the business.
  - Set goals so that everyone knows if success is achieved.
- Setting production databases in noarchivelog mode and then relying on exports or cold backups.
  - No way to recover lost data if crash occurs after the backups.
  - Recommend turning archivelog mode ON and use RMAN or hot backups instead.
  - Need to **validate, test & regularly review backup & DR plans**
- Installing Oracle Enterprise Edition downloaded from technet.oracle.com and not buying any Oracle Support can be a

Jerry's Advice to George...

- In George's system you might find alert log, trace files, or reports from end users of errors that have been occurring for a long time but were not addressed. Then the problem either grew into something larger or due to changes in the business these issues grow into something much larger.
- Cleanup is not occurring on destination directories (bdump, cdump, udump, adump).
- Leaves temp to grow a ridiculously huge size. This just trains the developers to write un-optimized queries.
- **NEVER go without a development environment**
- **George also never tests anything.** How times have we been called in to fix in a production environment that would have been caught and corrected BEFORE they were a production problem if only some basic testing had been done?

Elaine Doesn't Work past 5 PM

Are you available when they need you?
When you are available, are you easy to deal with?

Elaine doesn't Work past 5 PM...

- Elaine doesn't understand the concept of the DBA.
- She doesn't understand the dedication needed
- If users have a problem after 5 PM, they wait until tomorrow for a solution.
- Elaine doesn't have the time each day to do the required maintenance tasks needed
- Elaine really wants to be (and should be) an Ad-hoc query user, but wants the salary of a DBA.
- She advises others to never accept a job as DBA, you can never escape!!! **She advises other DBAs to never give out home#/pager/cell number to developers...**

Jerry's Advice to Elaine...

- Don't set the max_dump_file_size set to the default (unlimited). Can take over file system if trace generated is bigger than the destination.
- Don't use Oracle 8i/9i settings in a 10g or 11g instance. Take advantage of new features and get rid of backward compatibility kernel settings.
- She has RAC, but has no redundancy on the private interconnect. Having the private interconnect compete with other public network traffic is a bad thing.
- NEVER apply a patch without testing in development first
- NEVER move code into production without test in development

Jerry's Advice to Elaine...

- ALWAYS document your environment, changes to the environment, and custom code – this makes life easier for those of us who have to support it
- Her datafiles are set to autoextend unlimited. Then you find out that the file systems they reside on are limited. An Oracle Error is triggered when Oracle tries to resize and can't find the space. Clients do this all the time because they rely on DBCA.
- Don't rely on the UNIX administrator's word that the filesystems underneath are I/O tuned. A good DBA should actually do the LUN recommendation, i.e. R1+0 for Oracle, R5 for backups, etc. Check File I/O for issues

Jerry is the Productive DBA

Can you stay up 63 hours? You won't need to in the future!

Jerry's Secret to Tuning; AWR Report and Grid Control

- Jerry is the Master Tuning Expert
- He Knows The Oracle
- He Leverages what he learned in Statspack
- He Learned what's new in AWR Report
- He applies his tuning skills to Grid Control
- He **Pro-actively tunes** to head off problems
- He Re-actively tunes when needed
- He lets Grid Control Tune for him
- **He's put his knowledge into Grid Control so that he can be more productive.**

Know the Oracle

"I admire risk takers. I like leaders – people who do things before they become fashionable or popular. I find that kind of integrity inspirational."

LAWRENCE J. ELLISON | *Chairman & Chief Executive Officer, 2003*

Oracle Firsts – *Innovation!*

1979 First commercial SQL relational database management system
1983 First 32-bit mode RDBMS
1984 First database with read consistency
1987 First client-server database
1994 First commercial and multilevel secure database evaluations
1995 First 64-bit mode RDBMS
1996 First to break the 30,000 TPC-C barrier
1997 First Web database
1998 First Database - Native **Java** Support; Breaks 100,000 TPC-C
1998 First Commercial RDBMS ported to **Linux**
2000 First database with **XML**
2001 First middle-tier database cache
2001 First RDBMS with **Real Application Clusters**
2004 First **True Grid Database**
2005 First **FREE Oracle Database** (10g Express Edition)
2006 First **Oracle Support for LINUX Offering**
2007 **Oracle 11g Released!**
2008 Oracle Exadata Server Announced (Oracle buys BEA)
2009 Oracle buys Sun – Java; MySQL; Solaris; Hardware; OpenOffice
2010 Oracle announces MySQL Cluster 7.1, Exadata, Exalogic

In 2007: Version 11g was Released

- The Focus has been Acquisitions and gaining **Market Share**
- Oracle 11g Database extends an already large lead
  - Easier to Manage the Database – Better Grid Control
  - Self Tuning through a variety of tools (Makes 1 person equal 10)
  - Better Security/Encryption & Recoverability via Flashback
  - Better Testing Tools (Real Application Testing)
- Andy Mendelsohn is still the database lead
- Releases of **Siebel, PeopleSoft, JDE** and Oracle12 Apps.
- New Oracle BI Suite & **Acquisition of Hyperion**

Tuning - Leverage ALL of your Knowledge

Do Developers think of this when they think of their DBA?

Tuning in General

- Both an Art and a Science – You make miracles!
- Exceptions often rule the day…Not a "one size fits all"
- Hardware & Architecture must be right for your application or it will be difficult to succeed.
- Enterprise Manager (also 3rd party products) are best for simple tuning and ongoing maintenance.
- V$/X$ are best for drilling deep into problems

- **Enterprise Manager 11g radically makes you better!**

Check Regularly

1. **Top 5 wait events**
2. **Load Profile**
3. Instance Efficiency Hit Ratios
4. **Wait Events**
5. **Latch Waits**
6. **Top SQL**
7. **Instance Activity**
8. **File I/O**
9. Memory Allocation
10. Undo

AWR – Load Profile

Report Summary

Cache Sizes

	Begin	End		
Buffer Cache:	10,240M	10,240M	Std Block Size:	8K
Shared Pool Size:	1,264M	1,264M	Log Buffer:	36,744K

Load Profile

	Per Second	Per Transaction
Redo size:	37,741,608.27	5,236,744.44
Logical reads:	239,964.89	33,295.74
Block changes:	137,275.83	19,047.37
Physical reads:	1.84	0.25
Physical writes:	4,708.71	653.35
User calls:	42.00	5.83
Parses:	24.05	3.34
Hard parses:	0.04	0.01
Sorts:	0.34	0.05
Logons:	0.71	0.10
Executes:	31.85	4.42
Transactions:	7.21	

% Blocks changed per Read:	57.21	Recursive Call %:	78.22
Rollback per transaction %:	25.00	Rows per Sort:	153.62

Done

AWR – Waits / Instance Efficiency

Instance Efficiency Percentages (Target 100%)

Buffer Nowait %:	99.31	Redo NoWait %:	99.99
Buffer Hit %:	100.00	In-memory Sort %:	100.00
Library Hit %:	99.94	Soft Parse %:	99.82
Execute to Parse %:	24.50	Latch Hit %:	94.65
Parse CPU to Parse Elapsd %:	91.87	% Non-Parse CPU:	99.96

Shared Pool Statistics

	Begin	End
Memory Usage %:	68.02	68.20
% SQL with executions>1:	81.94	81.36
% Memory for SQL w/exec>1:	77.24	74.72

Top 5 Timed Events

Event	Waits	Time(s)	Avg Wait(ms)	% Total Call Time	Wait Class
log buffer space	52,521	8,851	169	35.8	Configuration
CPU time		7,636		30.9	
log file sync	7,362	5,122	696	20.7	Commit
buffer busy waits	1,564,508	2,145	1	8.7	Concurrency
log file sequential read	35,171	701	20	2.8	System I/O

Done

Statspack (old/free way) – Top 5 Wait Events

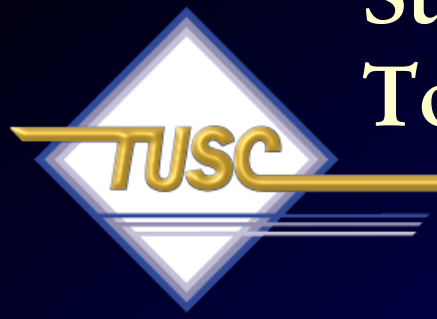

Top 5 Timed Events
~~~~~~~~~~~~~~~~~~

Event	Waits	Time (s)	% Total Ela Time
db file sequential read	399,394,399	2,562,115	52.26
CPU time		960,825	19.60
buffer busy waits	122,302,412	540,757	11.03
PL/SQL lock timer	4,077	243,056	4.96
log file switch (checkpoint incomplete)	188,701	187,648	3.83
~~~~~~~~~~~~~~~~~~

Top Wait Events
Things to look for...

Wait Problem	Potential Fix
Sequential Read I/O; Don't indexes.	Indicates many index reads – tune the code (especially joins); Faster over index or overuse
Scattered Read tune Faster I/O	Indicates many full table scans–index, the code; cache small tables;
Free Buffer to	Increase the DB_CACHE_SIZE; shorten the checkpoint; tune the code get less dirty blocks, faster I/O, use multiple DBWR's.
Buffer Busy	Segment Header – Add freelists (if

Statspack - Top Wait Events Things to look for...

Wait Problem	Potential Fix
Buffer Busy potentially	Data Block – Separate 'hot' data; use reverse key indexes; fix queries to reduce the blocks popularity, use smaller blocks, I/O, Increase initrans and/or maxtrans (this one's debatable) Reduce records per block.
Buffer Busy	Undo Header – Add segments or increase size of segment area (auto undo)
Buffer Busy	Undo block – Commit more (not too

Statspack - Top Wait Events Things to look for...

Wait Problem	Potential Fix
Enqueue - ST	Use LMT's or pre-allocate large extents
Enqueue - HW	Pre-allocate extents above HW (high water mark.)
Enqueue – TX (TX4) index. Fix	Increase initrans and/or maxtrans on (transaction) the table or locking issues if TX6. Bitmap (TX4) & Duplicates in Index (TX4).
Enqueue - TM	Index foreign keys; Check application

Why INITRANS Matter!

Transactions Moving through Oracle: *ITL & Undo Blocks*

User 1 – Updates Row# 1&2
User 2 updates Row 3

(There are also *In Memory Updates* (IMU) in 11g)

- *User1 updates a row with an insert/update/delete – an ITL is opened and xid tracks it in the data block.*
- *The xid ties to the UNDO header block which ties to the UNDO data block for undo.*
- *If user2 wants to query the row, they create a clone and rollback the transaction going to the undo header and undo block.*
- *If user3 wants to update same row (they wait). If user 3 wants to update different row then they open a second*

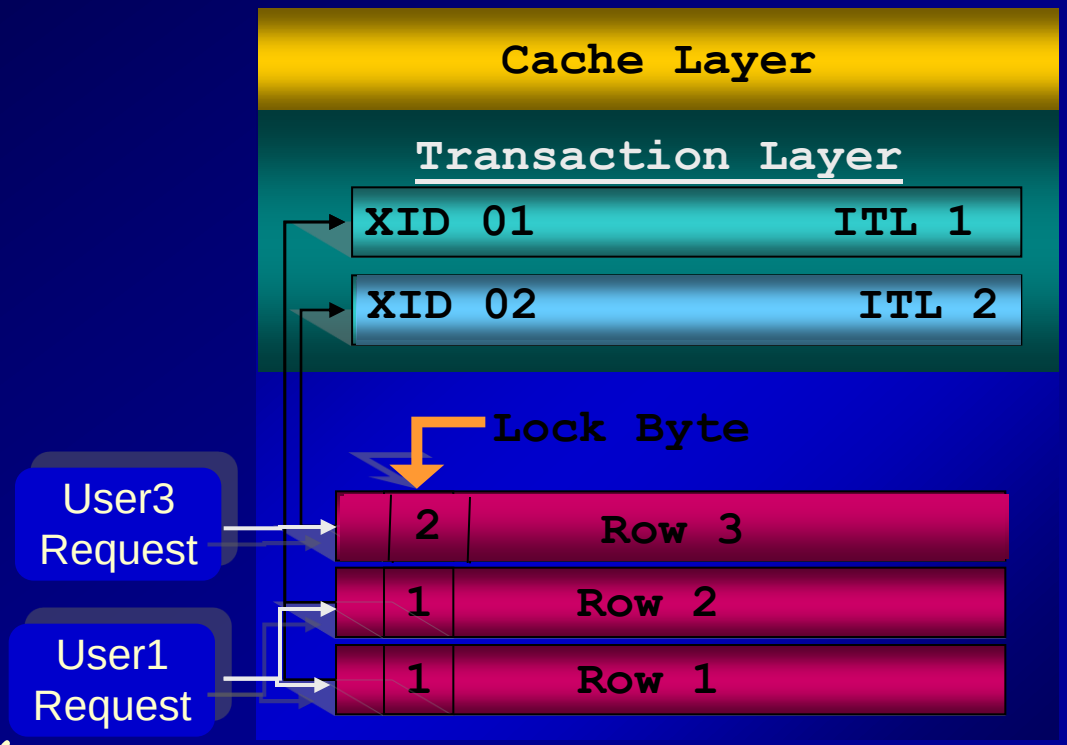

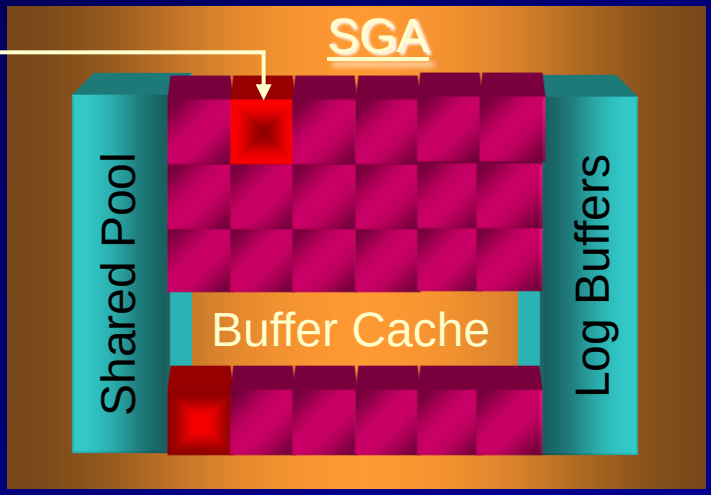

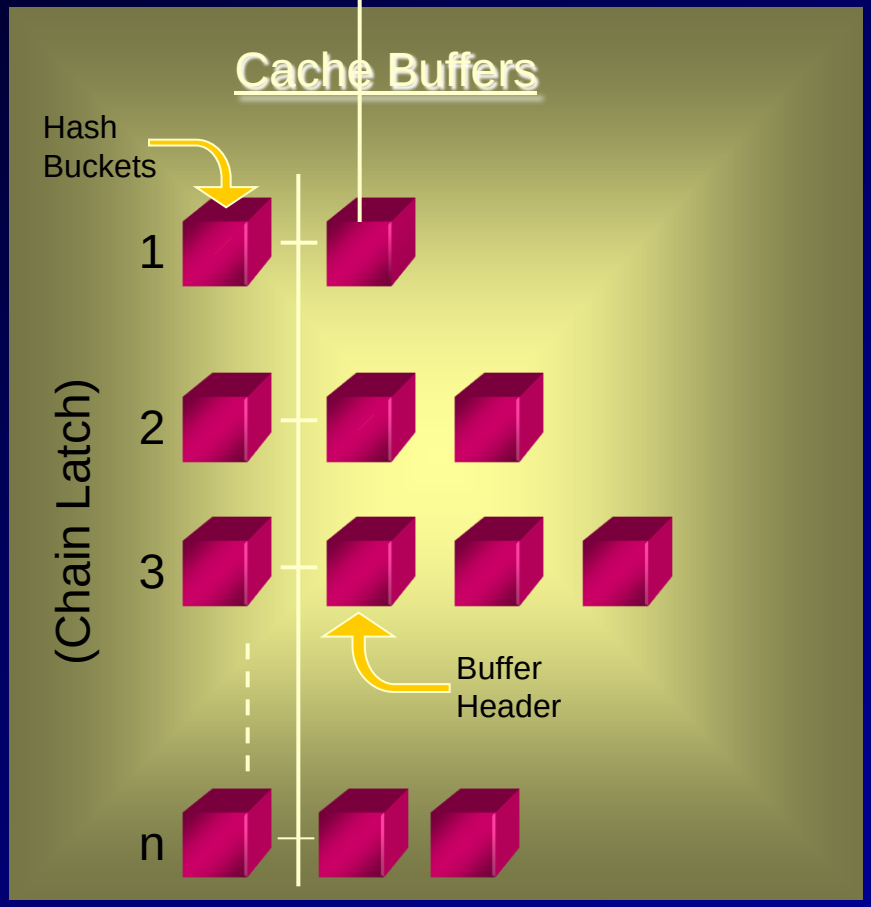

Only ONE block on the Hash Chain!

header
(So far it's clean and only 1 copy)

```
select     lrba_seq, state, dbarfil, dbablk, tch, flag, hscn_bas,cr_scn_bas,
    decode(bitand(flag,1), 0, 'N', 'Y') dirty, /* Dirty bit */
    decode(bitand(flag,16), 0, 'N', 'Y') temp,          /* temporary bit */
    decode(bitand(flag,1536), 0, 'N', 'Y') ping, /* ping (to shared or null) bit */
    decode(bitand(flag,16384), 0, 'N', 'Y') stale,      /* stale bit */
    decode(bitand(flag,65536), 0, 'N', 'Y') direct,     /* direct access bit */
    decode(bitand(flag,1048576), 0, 'N', 'Y') new       /* new bit */
from       x$bh
where      dablk = 56650
order by dablk;
```

```
  LRBA_SEQ       STATE     DBARFIL      DBABLK         TCH        FLAG    HSCN_BAS
---------- ---------- ---------- ---------- ---------- ---------- ----------
CR_SCN_BAS D T P S D N
---------- - - - - - -
         0          1          1      56650          0   35659776 4294967295
           0 N N N N N N
```

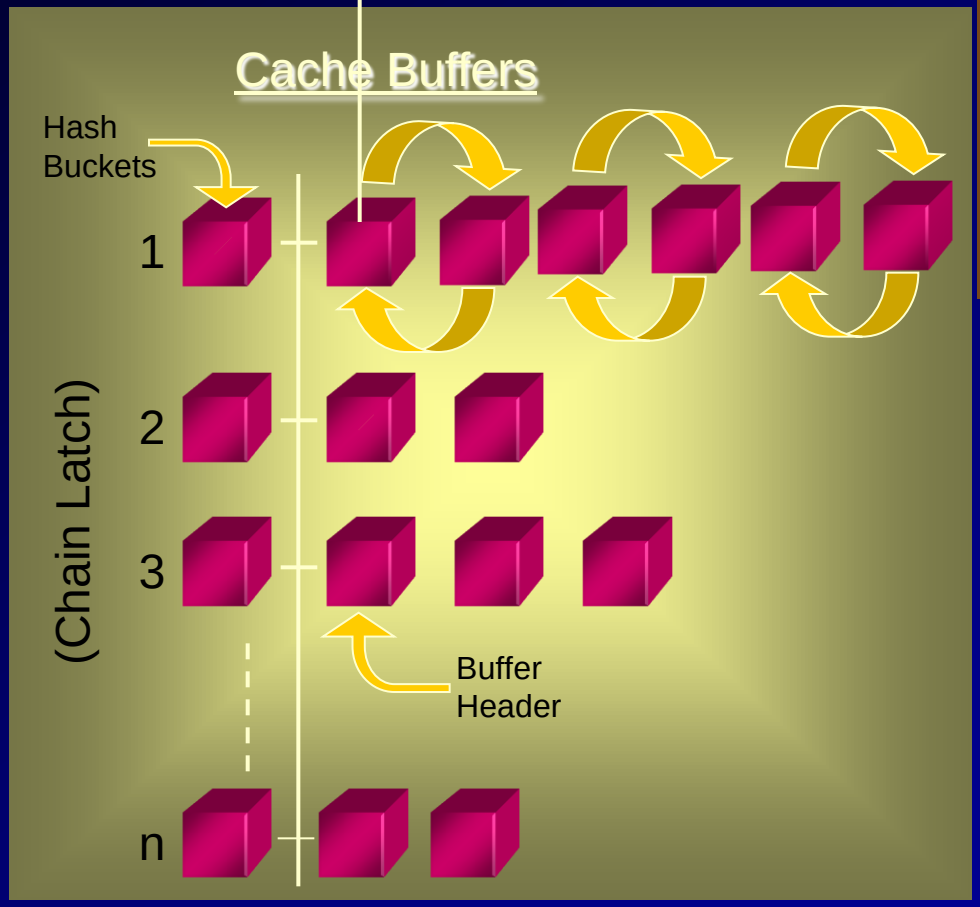

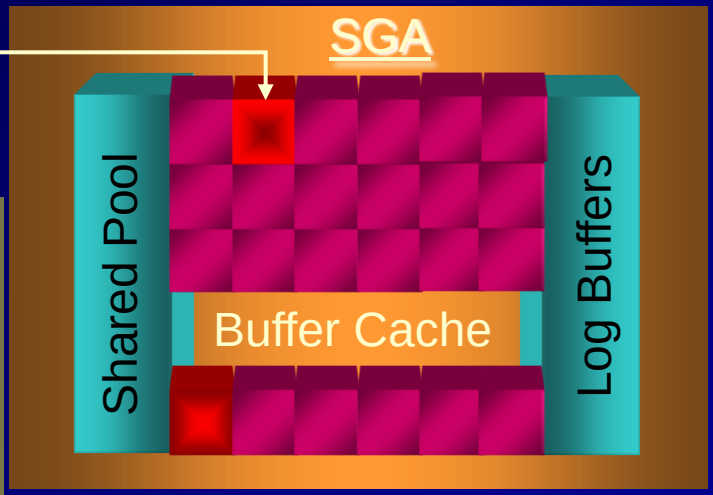

Hash Chain is now SIX long! Five CR and the one

x$bh up to the max of 6 versions of block


```
LRBA_SEQ      STATE     DBARFIL     DBABLK          TCH         FLAG    HSCN_BAS
---------- ---------- ---------- ---------- ---------- ---------- ----------
CR_SCN_BAS D T P S D N
---------- - - - - - -
         0          3          1      56650          1     524416          0
   4350120 N N N N N N

         0          3          1      56650          1     524416          0
   4350105 N N N N N N

       365          1          1      56650          7   33562633    4350121
         0 Y N N N N N

         0          3          1      56650          1     524416          0
   4350103 N N N N N N

         0          3          1      56650          1     524416          0
   4350089 N N N N N N

         0          3          1      56650          1     524288          0
   4350087 N N N N N N
```

Why only 6 versions of a Block?

```
select    a.ksppinm, b.ksppstvl, b.ksppstdf, a.ksppdesc
from      x$ksppi a, x$ksppcv b
where     a.indx = b.indx
and       substr(ksppinm,1,1) = '_'
and       ksppinm like  '%&1%'
order by            ksppinm;
```

```
KSPPINM
--------------------------------------------------------------------------------
KSPPSTVL
--------------------------------------------------------------------------------
KSPPSTDF
---------
KSPPDESC
--------------------------------------------------------------------------------
_db_block_max_cr_dba
6
TRUE
Maximum Allowed Number of CR buffers per dba
```

AWR – ITL Issues

Segments by ITL Waits

- % of Capture shows % of ITL waits for each top segment compared
- with total ITL waits for all segments captured by the Snapshot

Owner	Tablespace Name	Object Name	Subobject Name	Obj. Type	ITL Waits	% of Capture
				INDEX PARTITION	126	32.06
				INDEX PARTITION	112	28.50
				INDEX PARTITION	66	16.79
				INDEX PARTITION	65	16.54
				INDEX PARTITION	12	3.05

What are you Waiting on?

Is this your Ad-Hoc Query User or Network Administrator?

Statspack – Top 25

- Tuning the top 25 buffer get and top 25 physical get queries has yielded system performance gains of anywhere from 5 percent to 5000 percent.
- The SQL section of the statspack report tells you which queries to potentially tune first.
- **The top 10 of your SQL statements should usually not be more than 10 percent of your buffer gets or disk reads.**

Statspack – Top SQL (Top 2 are 5T & 3T of reads!!)

```
  Buffer Gets    Executions  Gets per Exec  %Total Time(s)  Time
   (s) Hash Value
--------------- ------------ -------------- ------ ------- --------
    627,226,570          117    5,360,910.9    4.7 9627.09 10367.04
Module: JDBC Thin Client
SELECT * FROM (select d1.tablespace_name, d1.owner, d1.segment_t
ype, d1.segment_name, d1.header_file, d1.extents, d1.bytes, d1.b
locks, d1.max_extents , d1.next_extent from sys.dba_segments d1
where d1.segment_type != 'CACHE'  and  tablespace_name not in (s
elect distinct tablespace_name from sys.dba_rollback_segs)  orde

    409,240,446      175,418        2,332.9    3.1 ####### 59430.83
Module:    ?  @sap10ci (TNS V1-V3)
SELECT "TABNAME" , "VARKEY" , "DATALN" , "VARDATA" FROM "KAPOL"
WHERE "TABNAME" = :A0 AND "VARKEY" LIKE :A1 ORDER BY "TABNAME" ,
 "VARKEY"
```

AWR – Top SQL
(Top 1 is 2T – Second one only 250M)

SQL ordered by Gets

- Resources reported for PL/SQL code includes the resources used by all SQL statements called by the code.
- Total Buffer Gets: 225,112,503
- Captured SQL account for 99.9% of Total

Buffer Gets	Executions	Gets per Exec	% Total	CPU Time (s)	Elapsed Time (s)	SQL Id	SQL Module	SQL Text
224,907,873	1,680	133,873.73	99.91	7568.17	19515.02			
31,779	6,881	4.62	0.01	2.61	2.61			
21,515	1,688	12.75	0.01	6.18	14.19			
19,827	1,688	11.75	0.01	2.71	10.71			
15,186	1,685	9.01	0.01	4.94	11.82			
13,501	1,685	8.01	0.01	2.84	9.72			
7,867	1	7,867.00	0.00	7.56	23.36			
4,783	1	4,783.00	0.00	0.76	1.09			
3,906	651	6.00	0.00	0.53	0.53			
2,640	646	4.09	0.00	0.62	1.21			

Back to SQL Statistics
Back to Top

SQL ordered by Reads

Statspack - Latch Waits

Latch Free – Latches are low-level queueing mechanisms (they're accurately referred to as mutually exclusion mechanisms) used to protect shared memory structures in the System Global Area (SGA).

- Latches are like locks on memory that are very quickly obtained and released.
- Latches are used to prevent concurrent access to a shared memory structure.
- If the latch is not available, a latch free miss is recorded.

Statspack - Latch Waits – fyi
Much better in 11g!!

Latch Free –

- Most latch problems are related to:
  - The failure to use bind variables (library cache latch)
  - Slow redo log disks or contention (log file sync)
  - Buffer cache contention issues (cache buffers lru chain)
  - Hot blocks in the buffer cache (cache buffers chains).
- There are also latch waits related to bugs; check Support for bug reports if you suspect this is the case (oracle.com/support).
- When latch miss ratios are greater than 0.5 percent, you should investigate the issue.
- **In memory updates have changed things for the better!**

Statspack - Latch Waits - fyi

Latch Activity for DB: ORA9I Instance: ora9i Snaps: 1 -2

Latch	Get Requests	Pct Get Miss	Avg Slps /Miss	Wait Time (s)	NoWait Requests	Pct NoWait Miss
KCL freelist latch	4,924	0.0			0	
cache buffer handles	968,992	0.0	0.0		0	
cache buffers chains	761,708,539	0.0	0.4		21,519,841	0.0
cache buffers lru chain	8,111,269	0.1	0.8		19,834,466	0.1
library cache	67,602,665	2.2	2.0		213,590	0.8
redo allocation	12,446,986	0.2	0.0		0	
redo copy	320	0.0			10,335,430	0.1
user lock	1,973	0.3	1.2		0	

Cursor Sharing - 8.1.6+

If v$sqlarea looks like this:

select empno from rich778 where empno =451572
select empno from rich778 where empno =451573
select empno from rich778 where empno =451574
select empno from rich778 where empno =451575
select empno from rich778 where empno =451576

Use cursor_sharing=force <similar> (sqlarea goes to this):

Instance Activity – AWR/Statspack

Statistic	Total	per Second	per Trans
branch node splits	7,162	0.1	0.0
consistent gets	12,931,850,777	152,858.8	3,969.5
current blocks converted for CR	75,709	0.9	0.0
db block changes	343,632,442	4,061.9	105.5
db block gets	390,323,754	4,613.8	119.8
hot buffers moved to head of LRU	197,262,394	2,331.7	60.6
leaf node 90-10 splits	26,429	0.3	0.0
leaf node splits	840,436	9.9	0.3
logons cumulative	21,369	0.3	0.0
physical reads	504,643,275	5,965.1	154.9
physical writes	49,724,268	587.8	15.3
session logical reads	13,322,170,917	157,472.5	4,089.4
sorts (disk)	4,132	0.1	0.0
sorts (memory)	7,938,085	93.8	2.4
sorts (rows)	906,207,041	10,711.7	278.2
table fetch continued row	25,506,365	301.5	7.8
table scans (long tables)	111	0.0	0.0
table scans (short tables)	1,543,085	18.2	0.5

Instance Activity Terminology... - fyi only

Statistic	Description
Session Logical Reads	All reads cached in memory. Includes both consistent gets and also the db block gets.
Consistent Gets	These are the reads of a block that are in the cache. They are NOT to be confused with consistent read (cr) version of a block in the buffer cache (usually the current version is read).
Db block gets	These are block gotten to be changed. MUST be the CURRENT block and not a CR block.
Db block changes	These are the db block gets (above) that were actually changed.
Physical Reads	Blocks not read from the cache. Either from disk, disk cache or O/S

File I/O

Tablespace

	Reads	Av Reads/s	Av Rd(ms)	Av Blks/Rd	Writes	Av Writes/s	Buffer Waits	Av Buf Wt(ms)
PSAPSTABI	14,441,749	171	7.9	1.0	521,275	6	1,234,608	6.2
PSAPVBAPD	13,639,443	161	6.2	1.7	10,057	0	2,672,470	4.2
PSAPEDII	11,992,418	142	5.3	1.0	83,757	1	4,115,714	4.4
PSAPEDID	10,617,042	125	8.1	1.0	64,866	1	3,728,009	6.4
PSAPROLL	998,328	12	13.2	1.0	8,321,252	98	285,060	65.7

- Reads should be below 14ms

AWR – File I/O

File IO Stats

- ordered by Tablespace, File

Tablespace	Filename	Reads	Av Reads/s	Av Rd (ms)	Av Blks/Rd	Writes	Av Writes/s	Buffer Waits	Av Buf Wt (ms)
		7	0	1.43	1.00	480	1	93	19.68
		6	0	0.00	1.00	34	0	34	16.76
		3	0	0.00	1.00	3	0	0	0.00
	.dbf	7	0	5.71	1.00	4	0	0	0.00
	.dbf	7	0	0.00	1.00	4	0	0	0.00
	.dbf	7	0	0.00	1.00	4	0	0	0.00
	.dbf	7	0	0.00	1.00	4	0	0	0.00
	.dbf	7	0	0.00	1.00	4	0	0	0.00
	.dbf	7	0	2.86	1.00	4	0	0	0.00
	.dbf	7	0	0.00	1.00	4	0	0	0.00
	.dbf	7	0	1.43	1.00	4	0	0	0.00
	.dbf	7	0	2.86	1.00	4	0	0	0.00
	.dbf	7	0	0.00	1.00	4	0	0	0.00
	.dbf	7	0	0.00	1.00	4	0	0	0.00
	.dbf	8	0	0.00	1.00	3	0	0	0.00
	.dbf	8	0	0.00	1.00	3	0	0	0.00
	.dbf	8	0	0.00	1.00	3	0	0	0.00
	.dbf	8	0	0.00	1.00	3	0	0	0.00

Done Internet

Grid Control – Find Problems FAST

*** TEST SCENARIO ***

Top Activity Shows Lots of Waiting Issues with Locking

Top Activity says BIG PROBLEMS

Almost 200 users are Active

We have some DML Issues

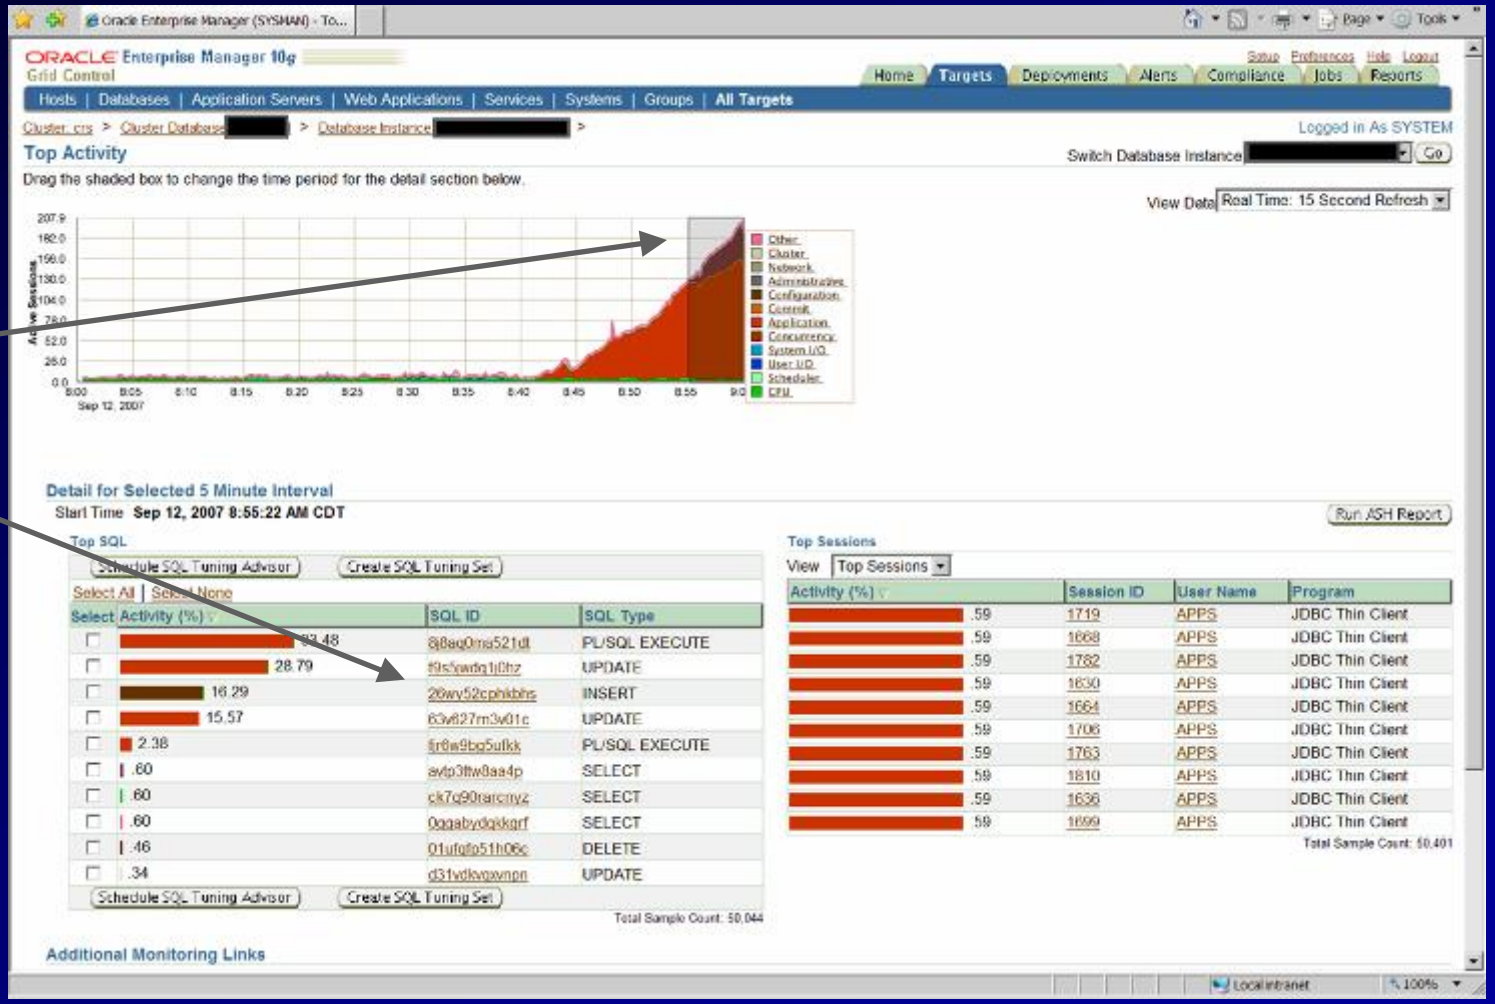

Top SQL #1 Details – Locking Issue

Tax Package shows row lock enqueue issue

MANY users are dividing the activity

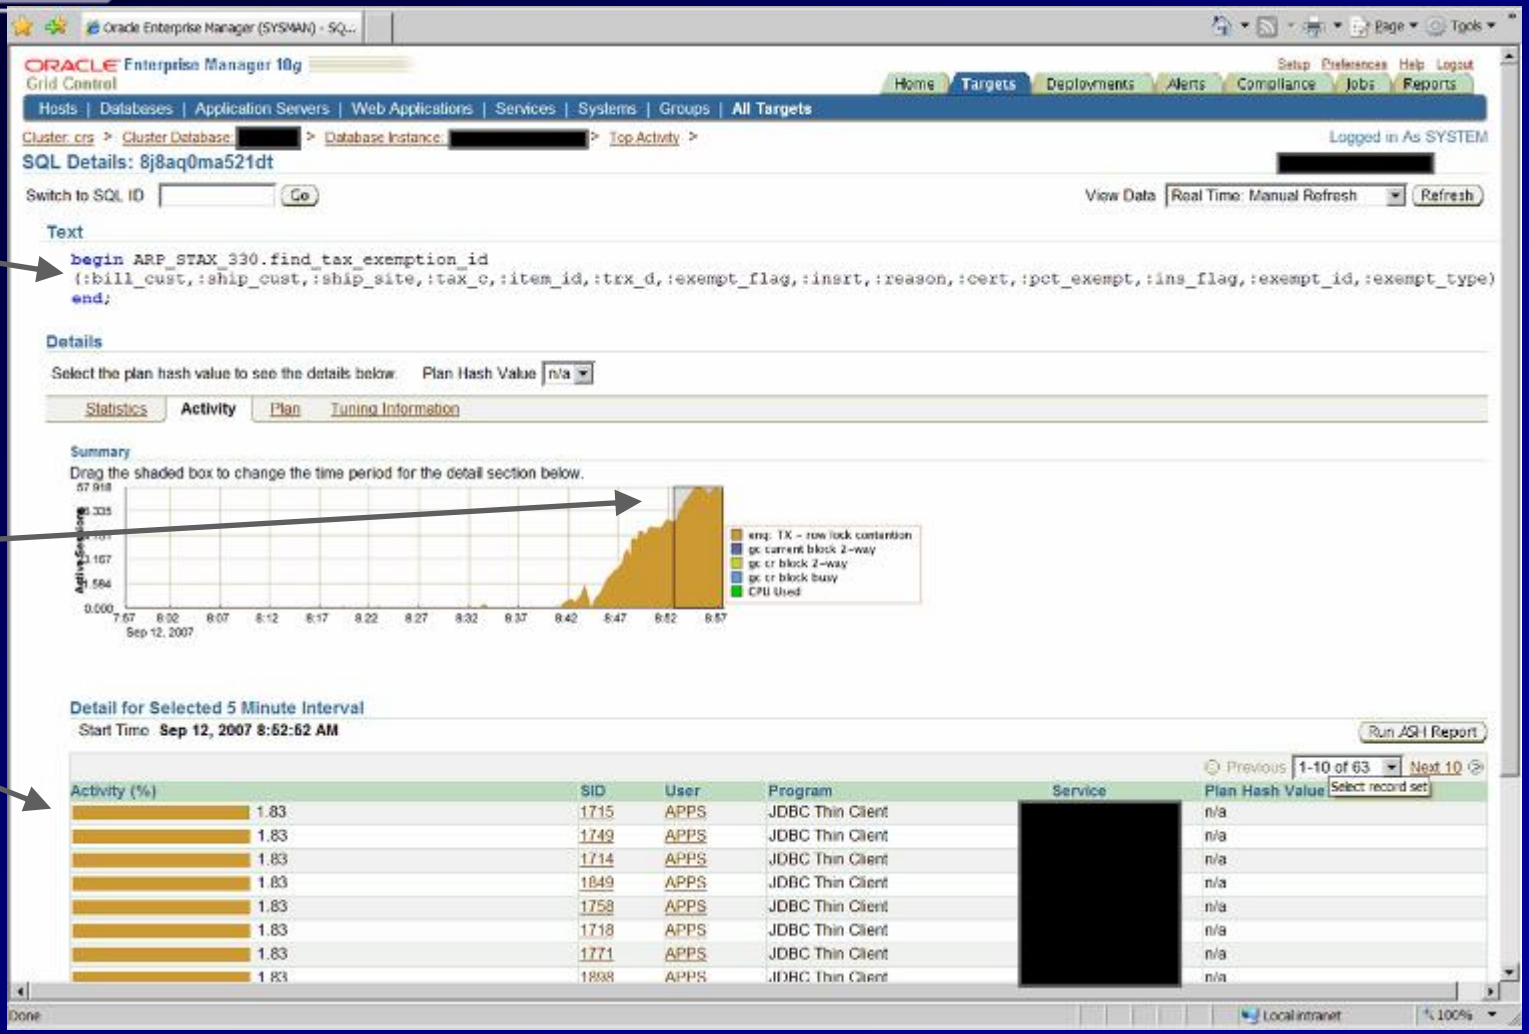

Histogram for Enqueues – Long Waits

LONG waits for the TX row lock

Top SQL #2 – Update Statement

The query in the causing the locks

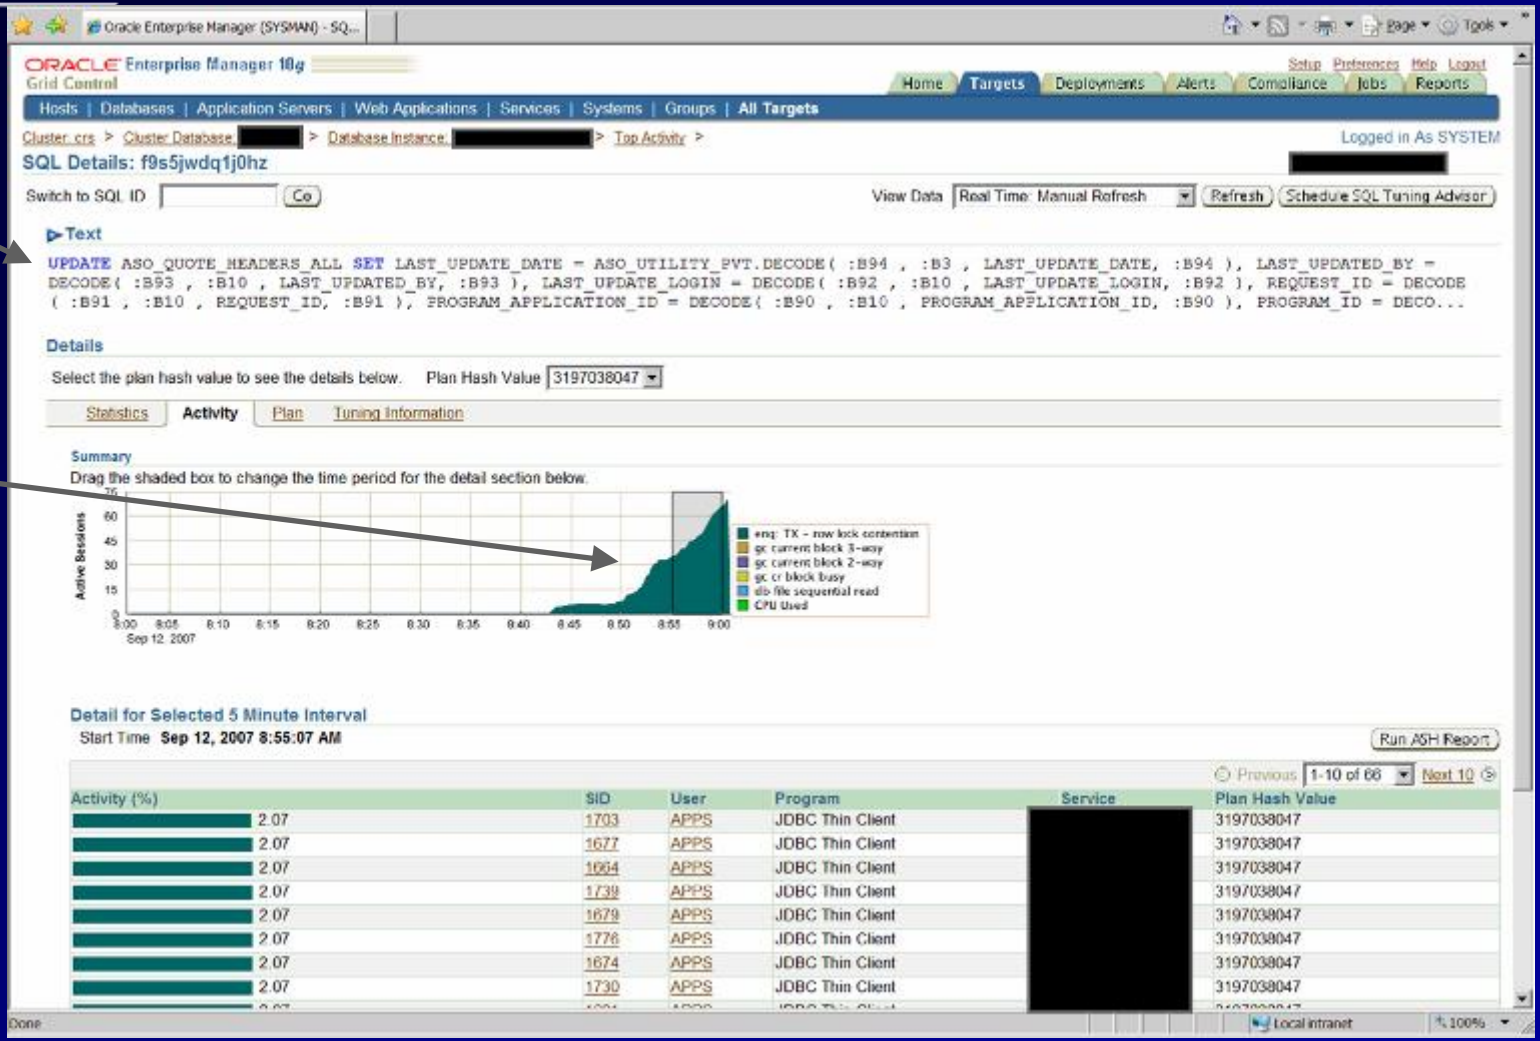

Top SQL #3 – Insert Statement

The application is causing BIG problems

Now there are over 400 active users

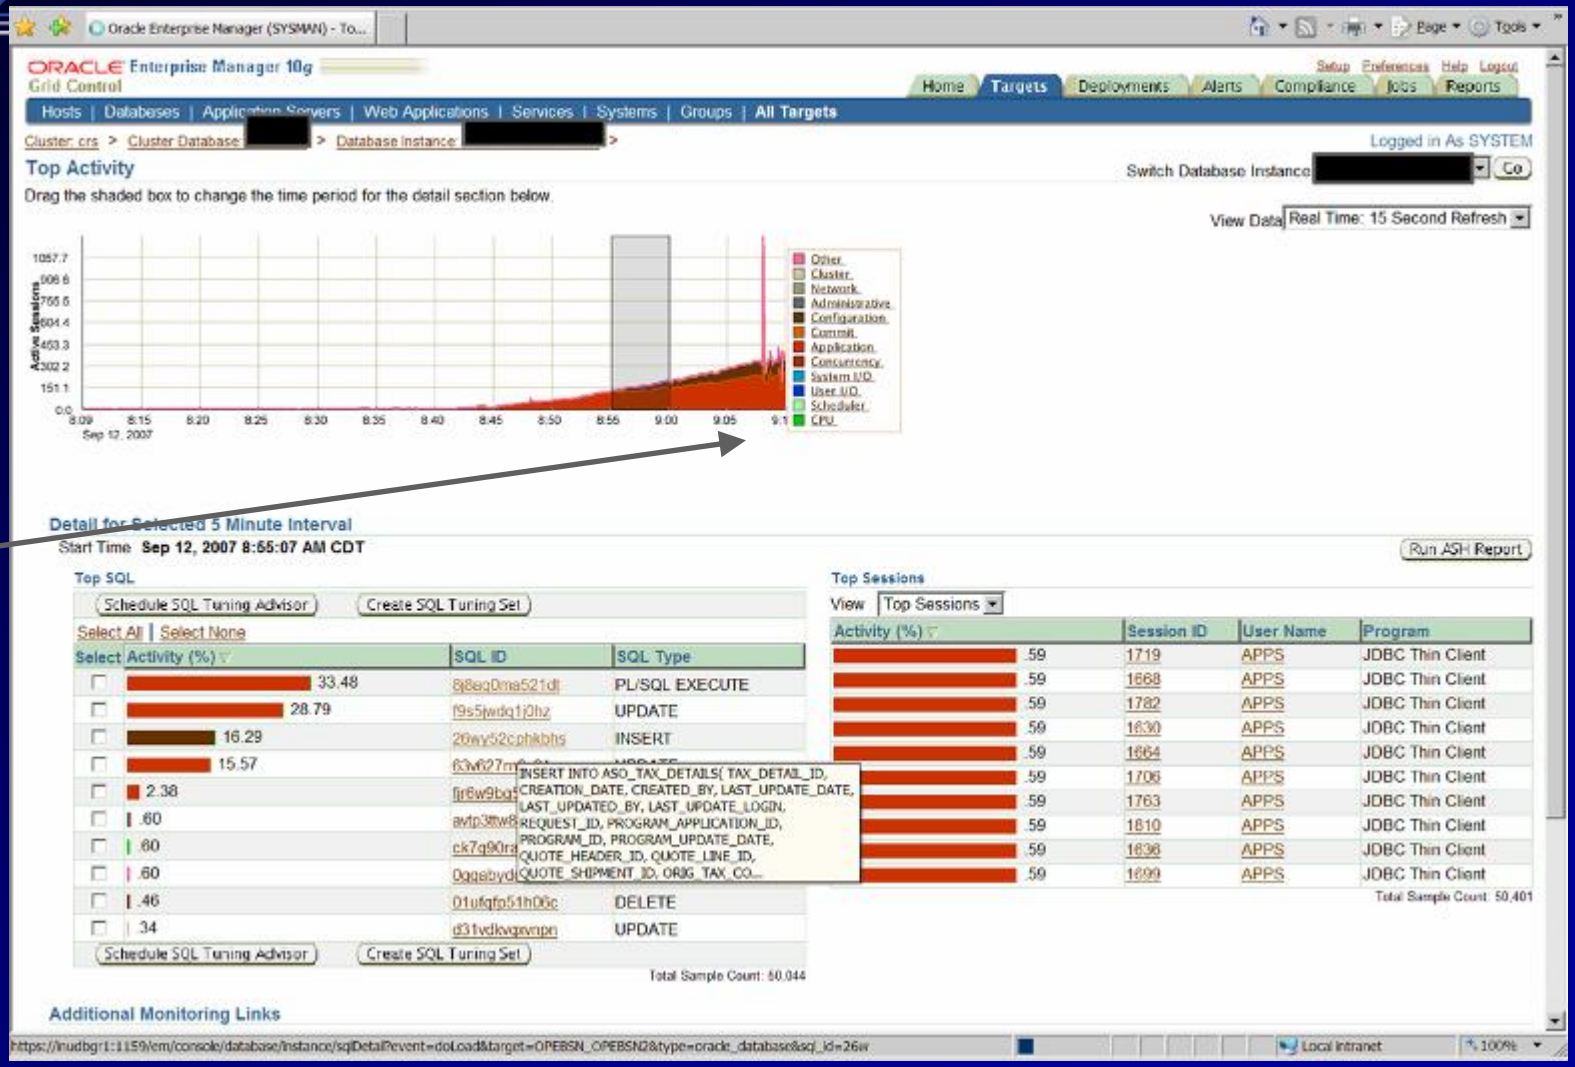

Top SQL #3 – Insert Statement Enqueue waits related to ITL allocations

The Insert into one of the TAX Tables

ITL issues

Some minor RAC gc issues

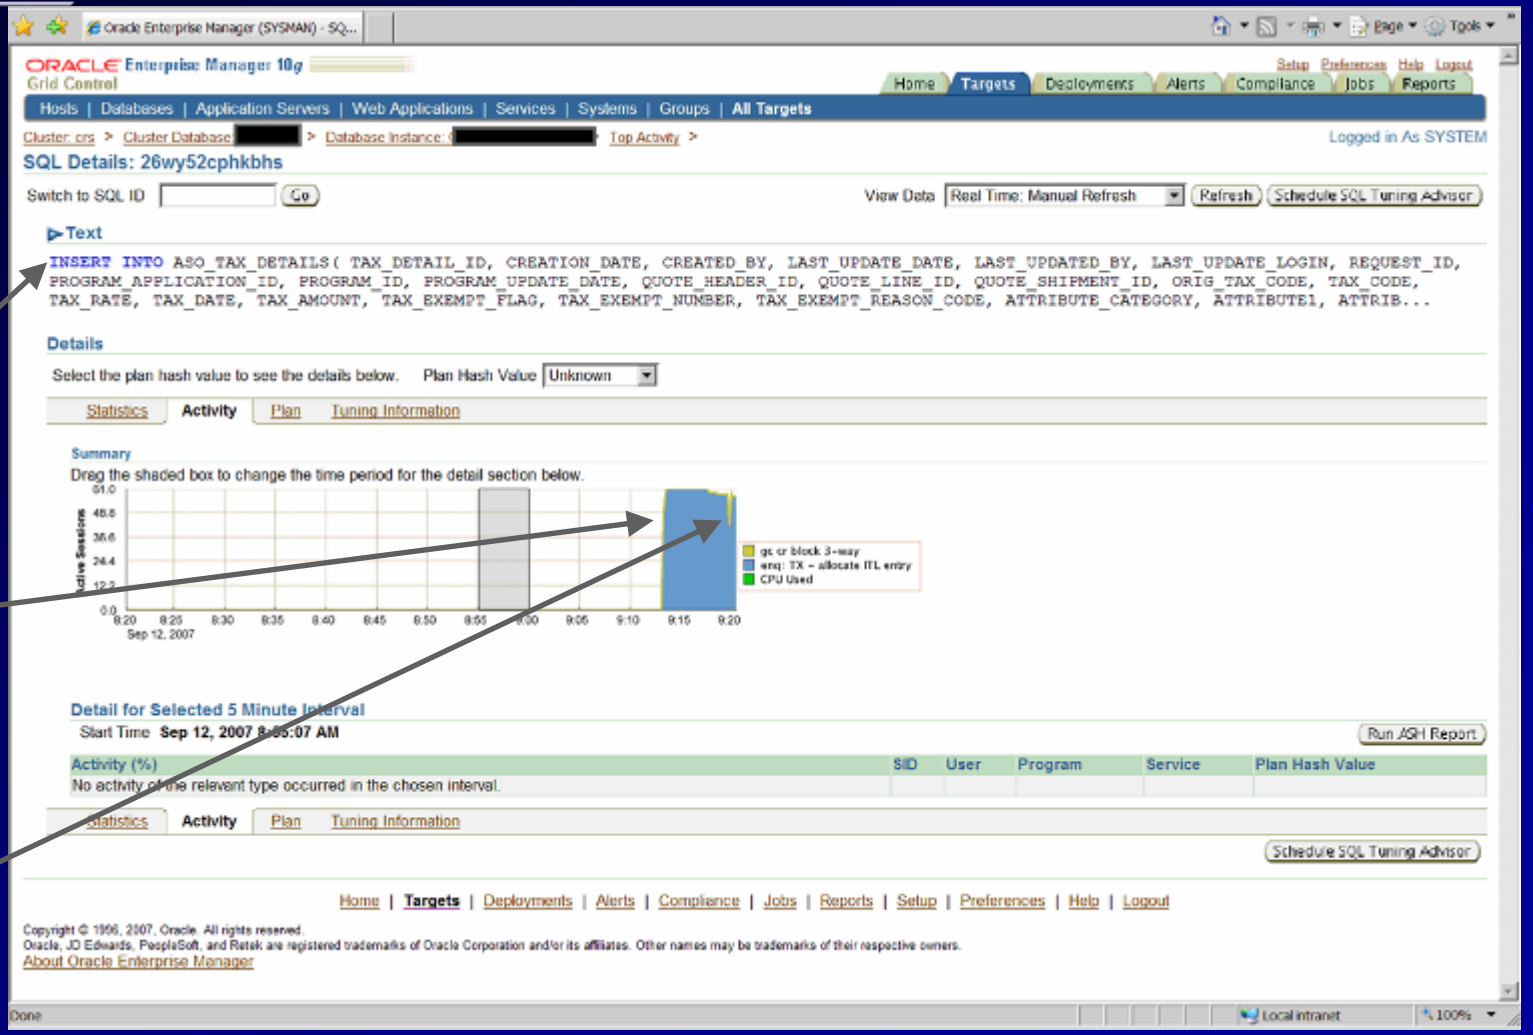

ADDM Enhancements
(Automatic Database Diagnostic Monitor)

ADDM enhancements

- Global ADDM so that Diagnostics are done across the entire cluster
- Emergency ADDM for use when database is hung
- On any granularity
  - Database Cluster
  - Database Instance
  - Specific Target (such as host, ASM...etc.)
- Over a specified time NOT tied to a pair of

ADDM Briefly

Specific Database Instance

We have 5 ADDM Findings

Check them Here

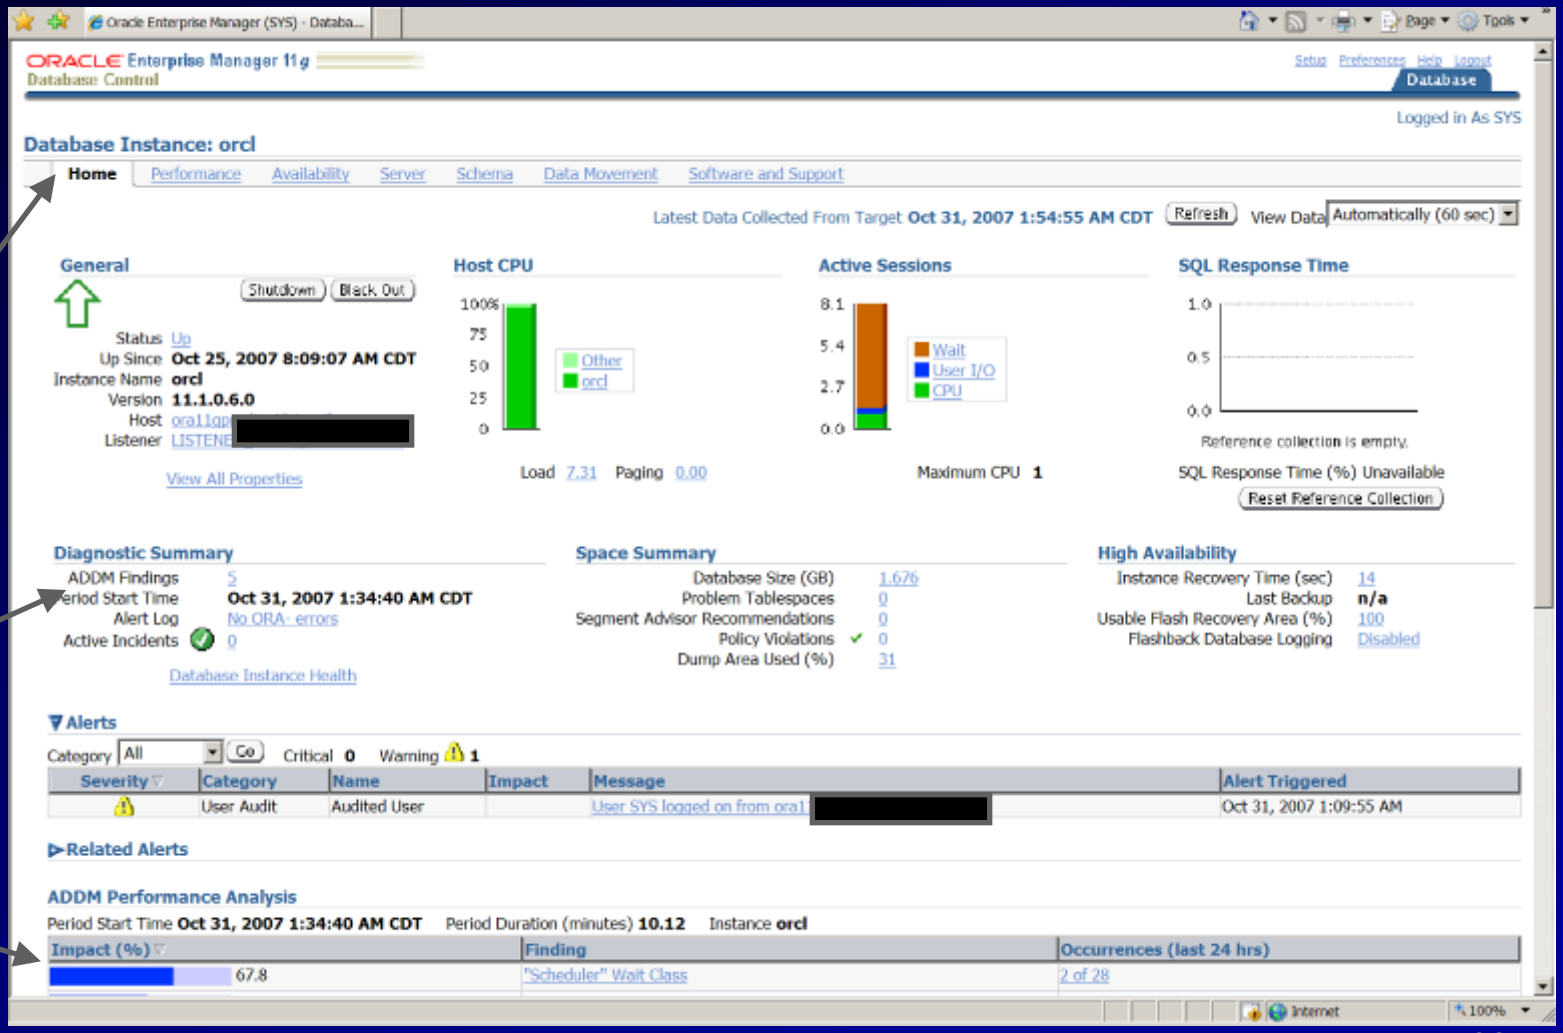

ADDM Briefly

Top ADDM Findings

Click a Single Timeframe

Let's Check the Hard Parse Issue

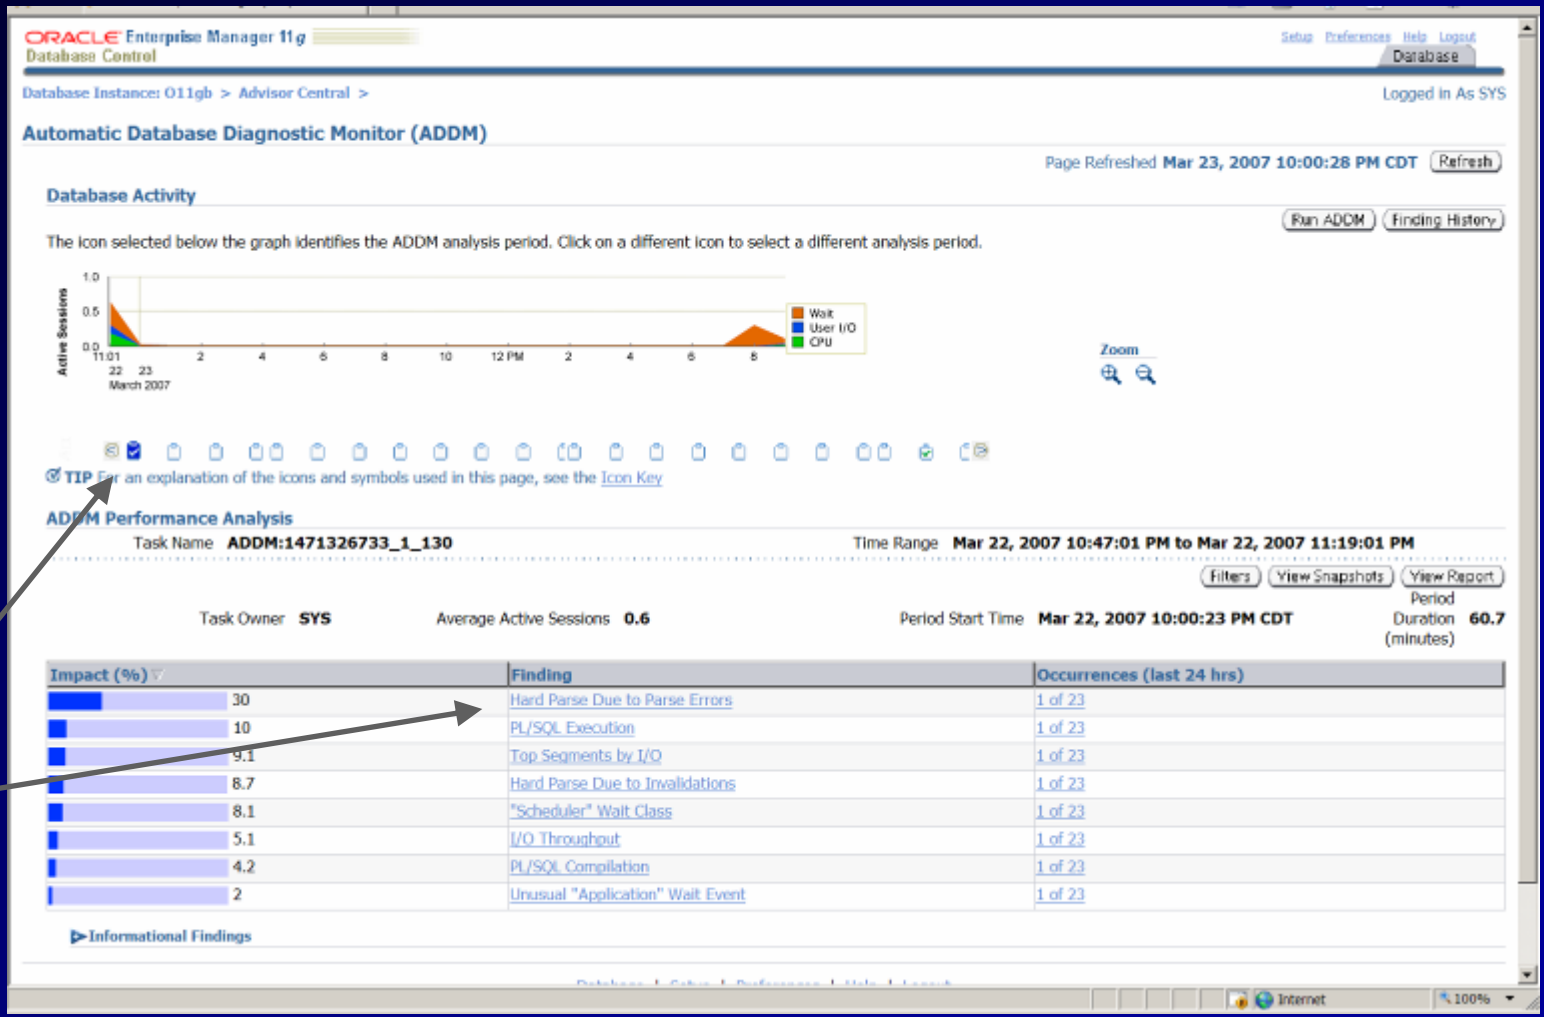

ADDM Briefly

Detailed Info & Findings

Add'l Info

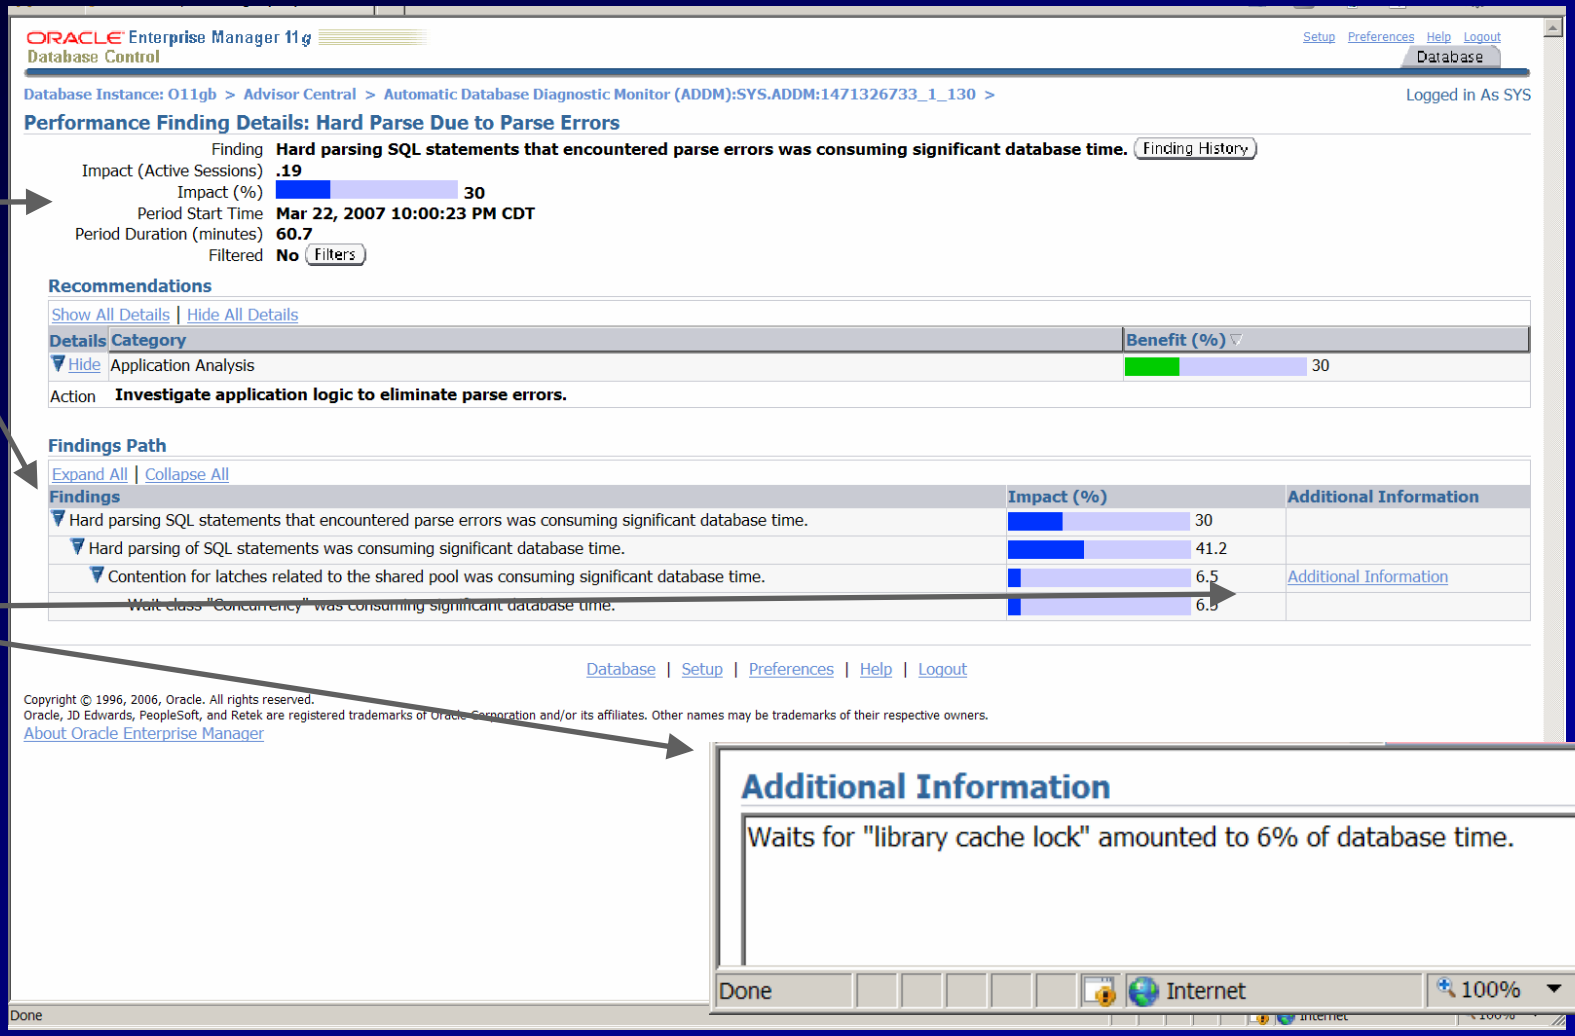

ADDM - Run NOW!

A Big Problem Occurs

Run ADDM NOW!

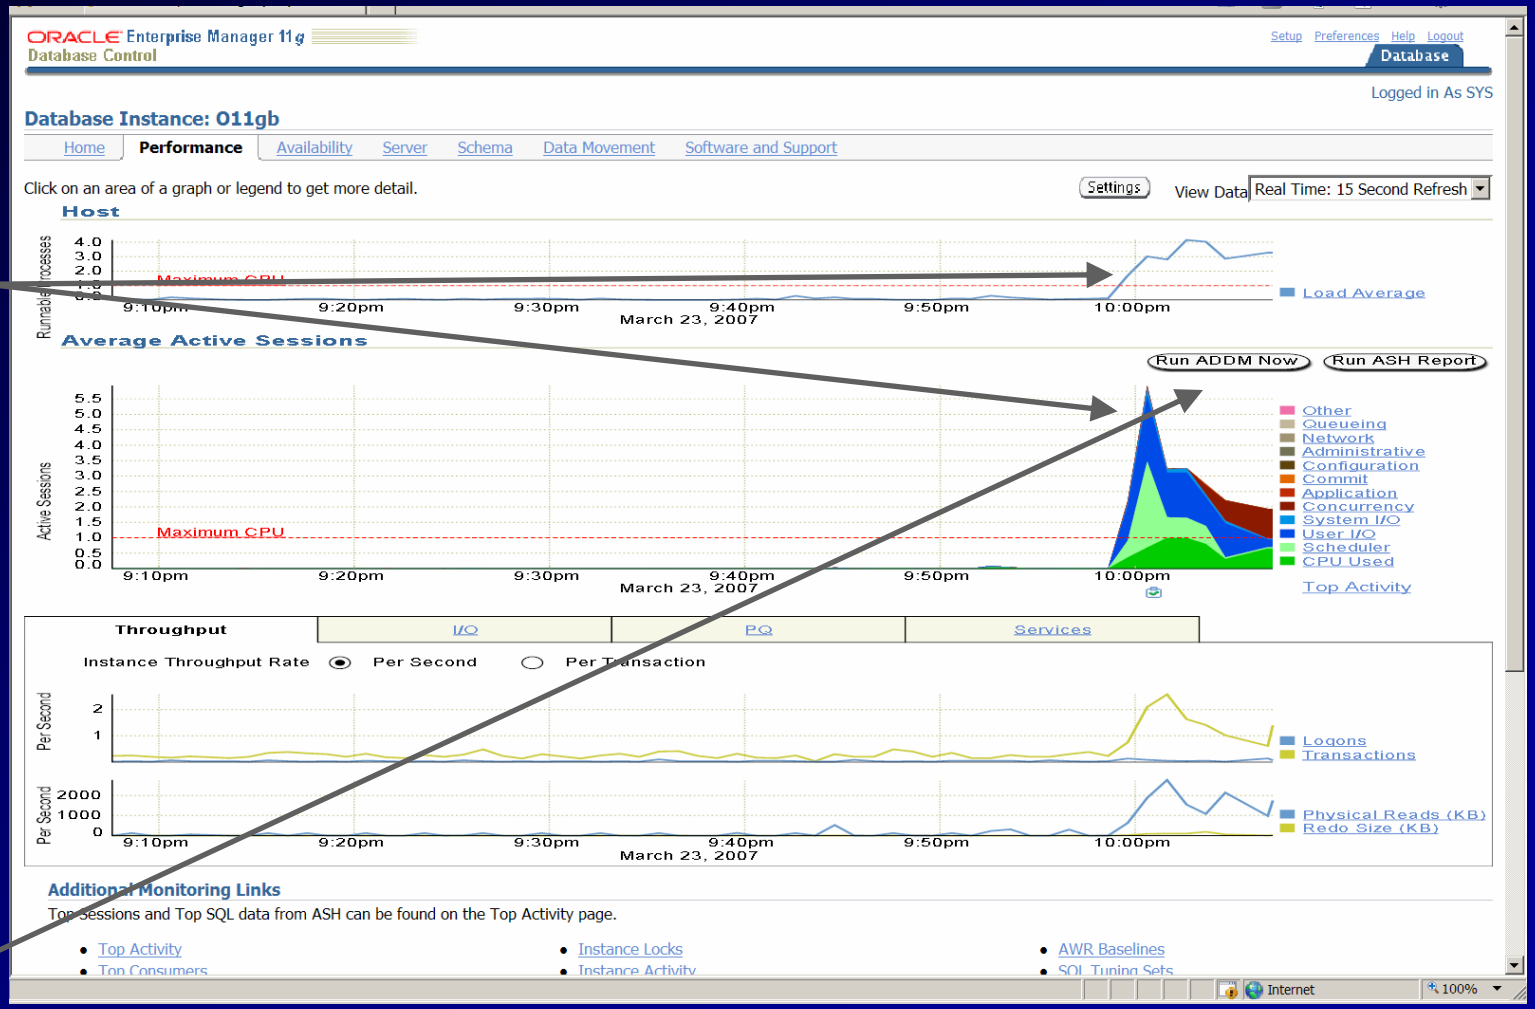

ADDM - Run NOW!

Are you Sure?

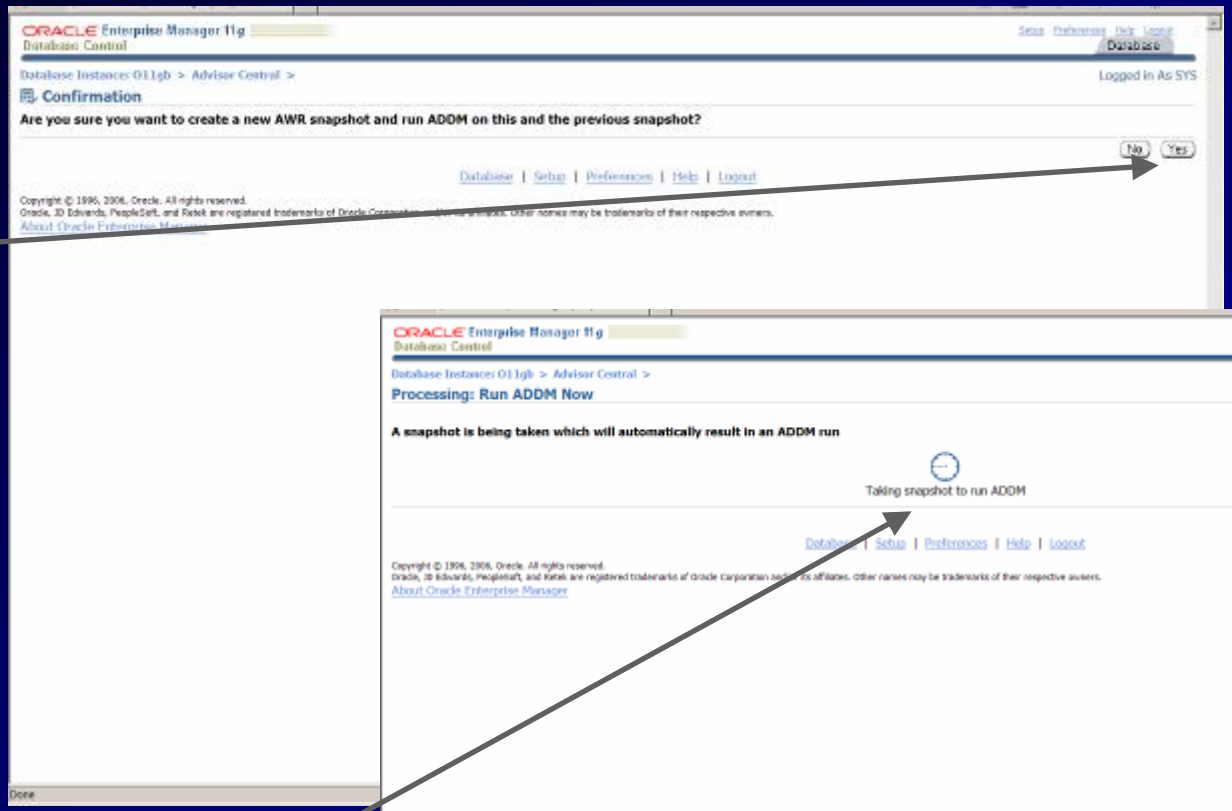

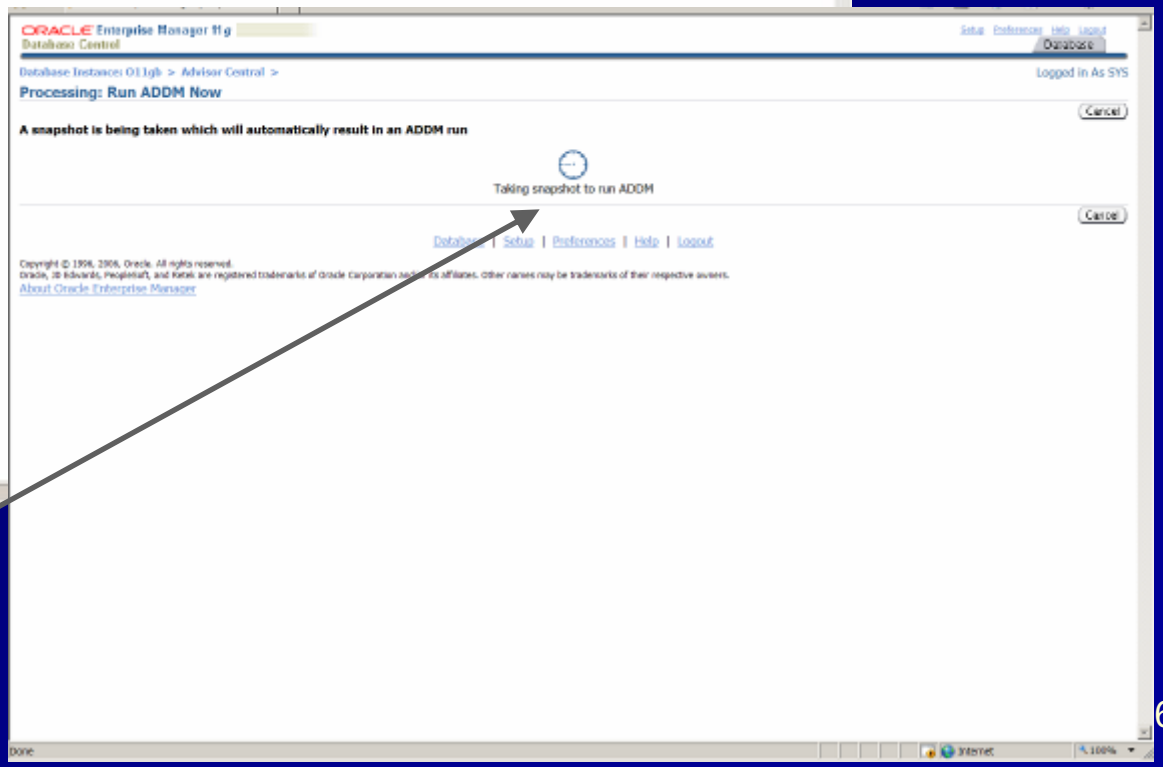

Running

ADDM – Run NOW!

Done.

CPU Issue

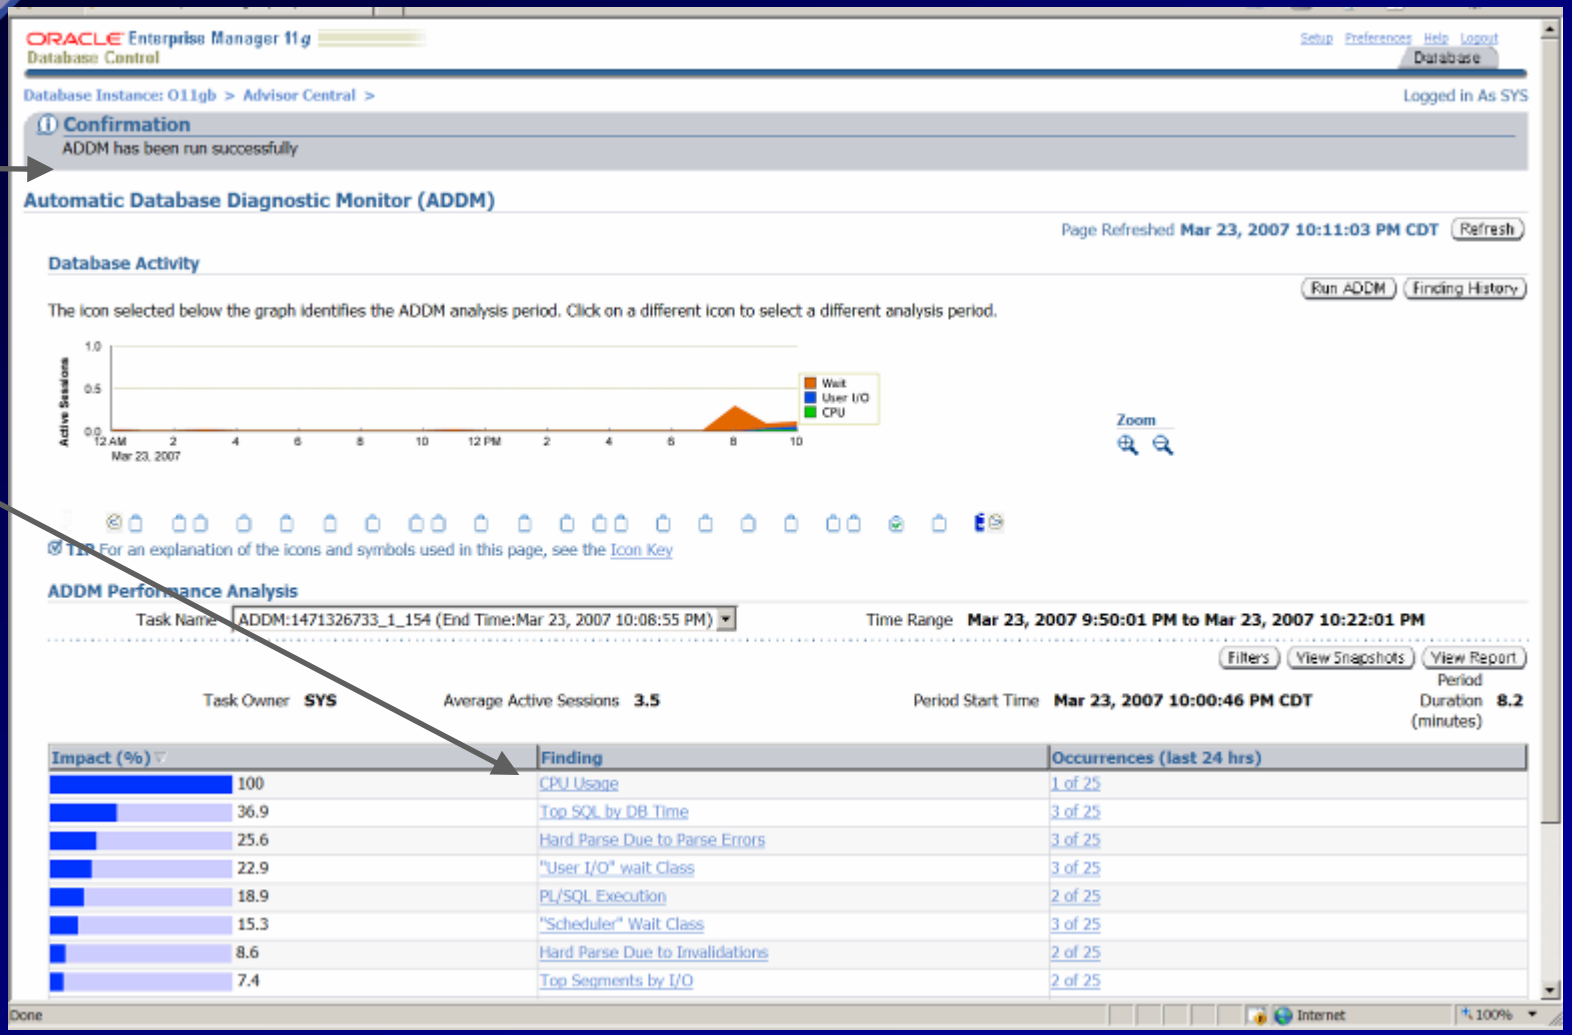

ADDM – Run NOW!

Detail on CPU Issue?

Suggested Fixes

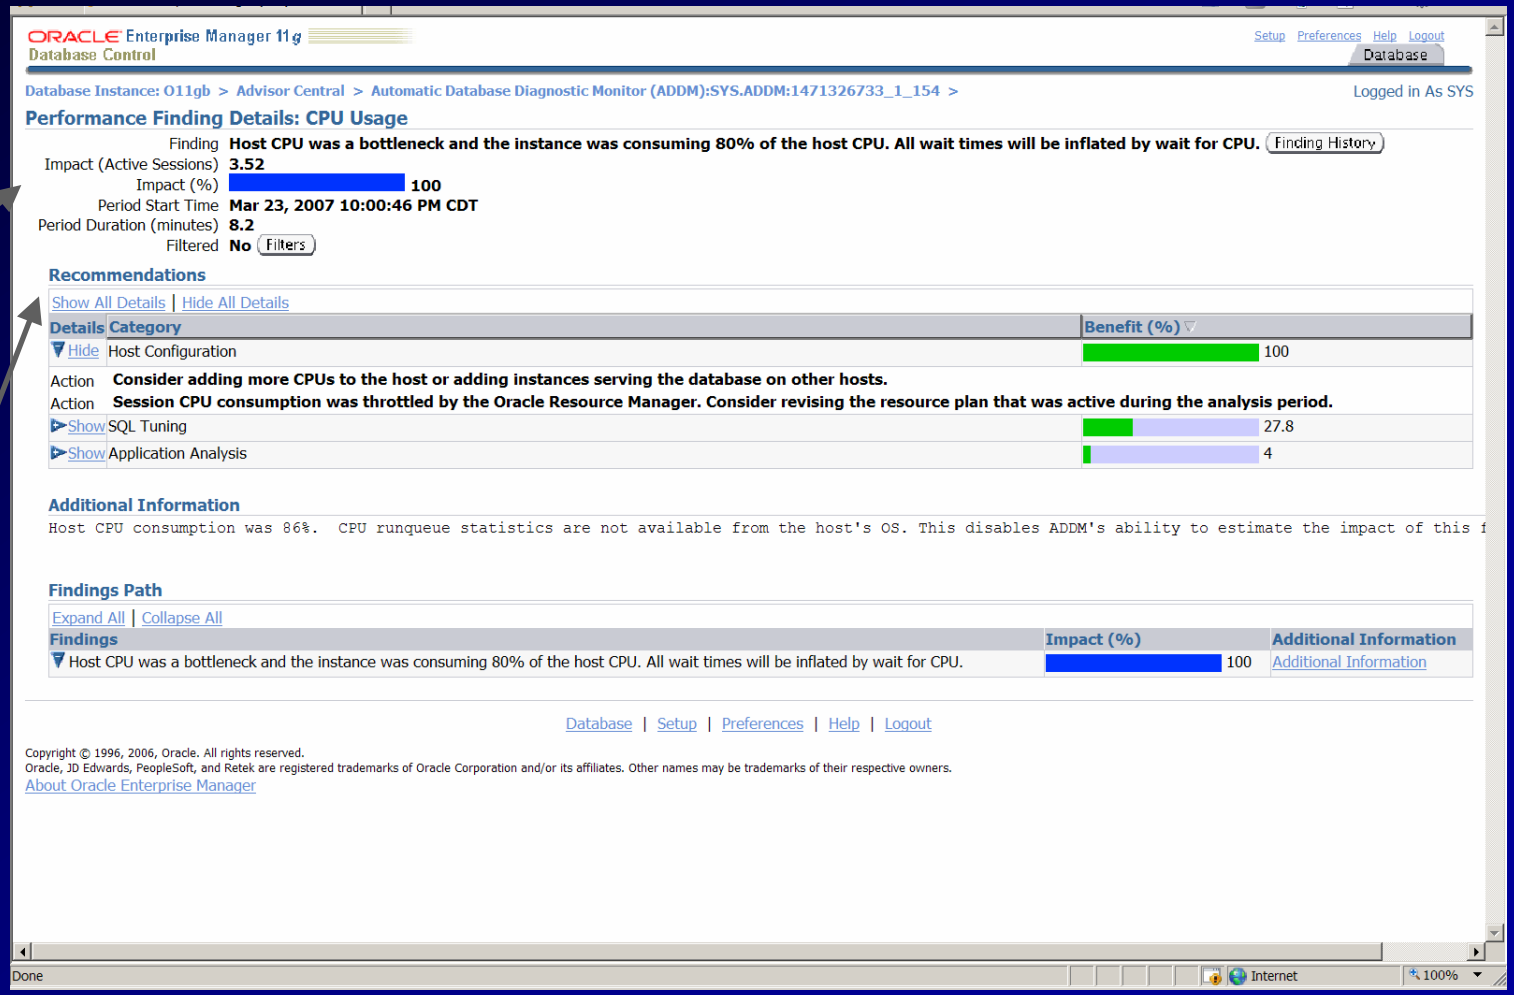

ADDM – Run NOW!

View The Report

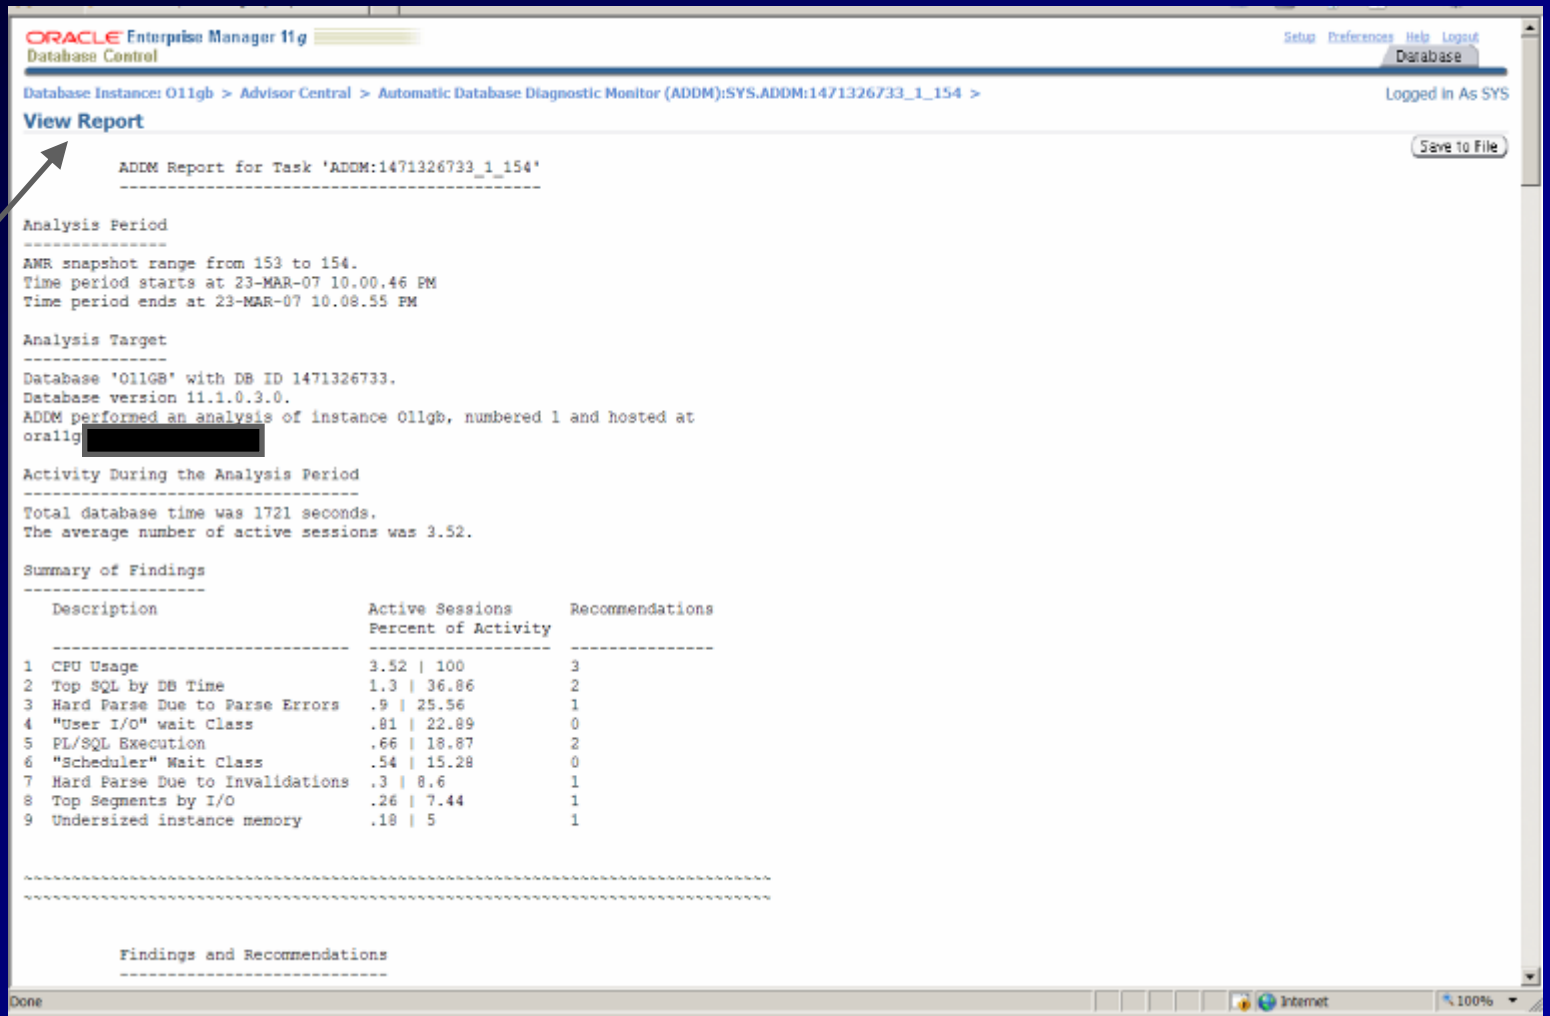

ADDM for RAC

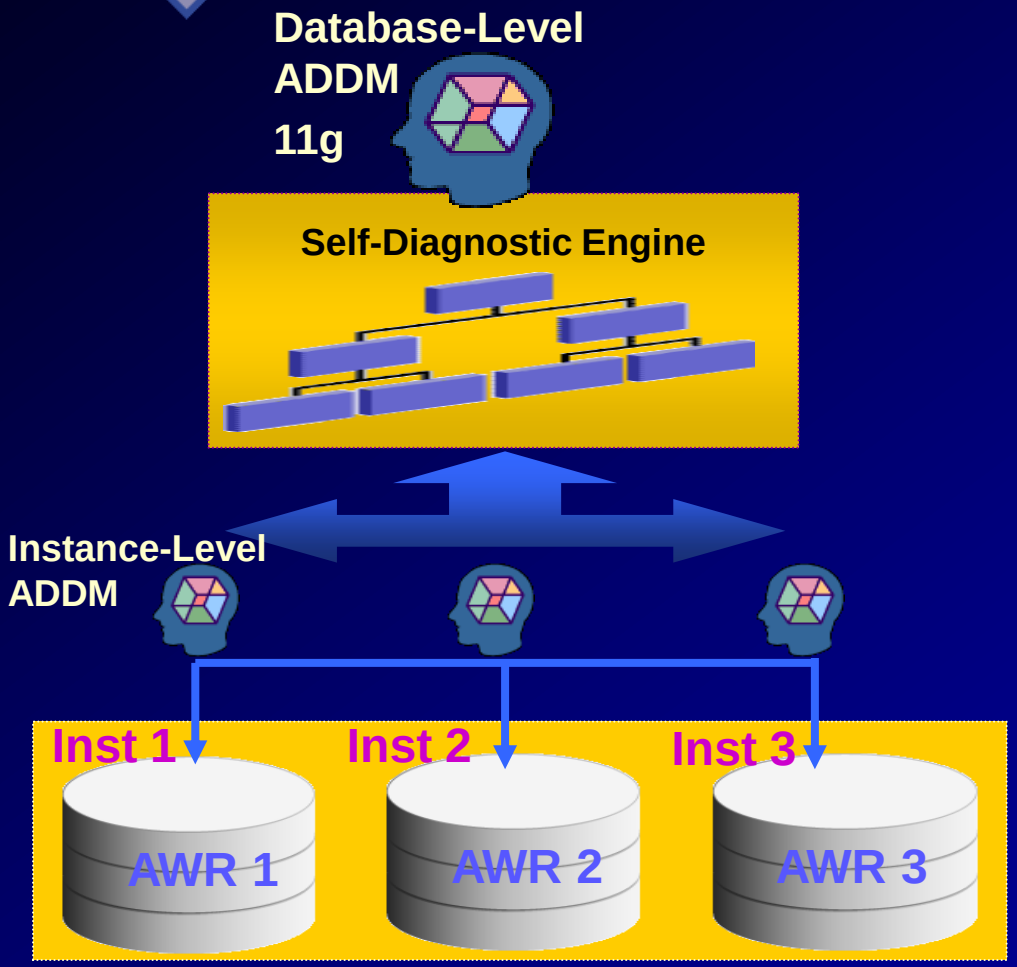

- Performance expert in a box
  - Now RAC specialist too!
- Identifies the most "Globally Significant" performance issues for the entire RAC database
- Database-wide and instance-level analysis
- Database-wide analysis of:
  - Global cache interconnect issues
  - Lock manager congestion issues
  - Global resource contention, e.g. IO bandwidth, hot blocks
  - Globally high-load SQL
  - Skew in instance response times
- Allows drill down to instances
- Runs proactively every hour when taking AWR Snapshots (default)

ADDM Considerations:

- CPU Bottlenecks
- Undersized Memory Structures – SGA / PGA
- I/O Capacity Issues
- High Load SQL statements
- High Load PL/SQL
- RAC specific issues – Global hot block/interconnect
- Application issues such as parsing, locks...etc.
- Concurrency (buffer busy) or hot object issues
- Configuration issues – Redo, Archive

Grid Control = Find Problems FAST

*** TEST SCENARIO ***

Go to other Tools to Verify & Advise:

ADDM

Hang Analysis

ASH

Go to ADDM to get Verify & Advise <5 minutes later>

ADDM sees the row contention for the Update

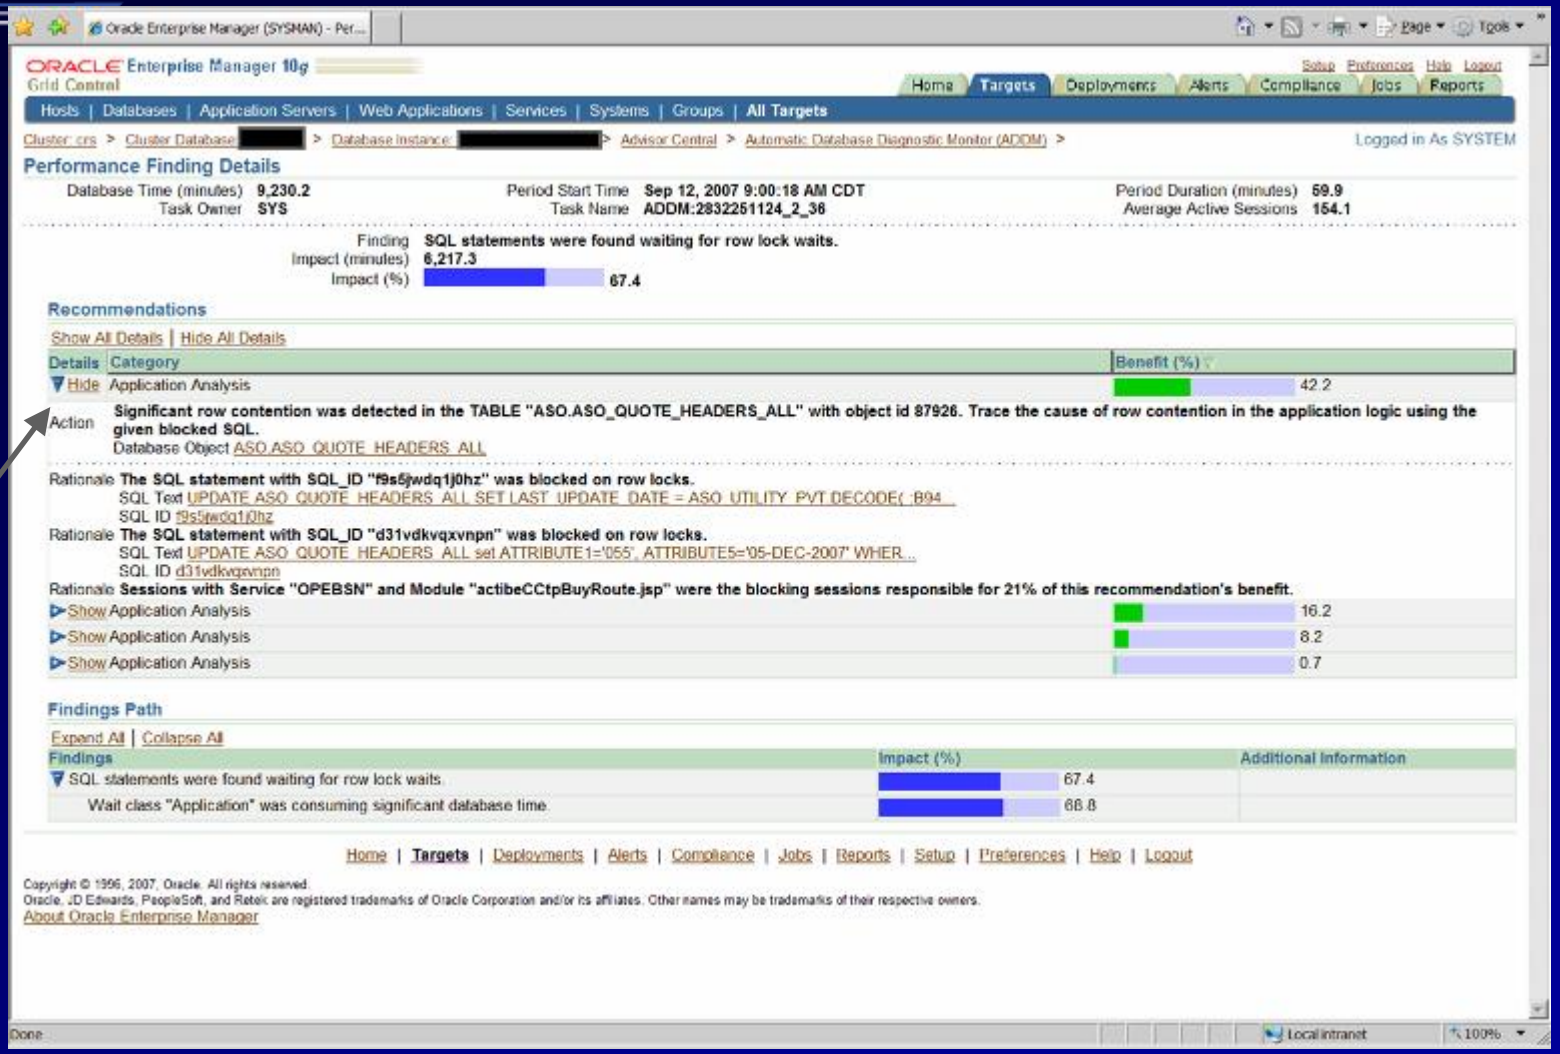

Grid Control Lessons Learned

Specific Update Statement for Tuning Results

Suggests to gather statistics

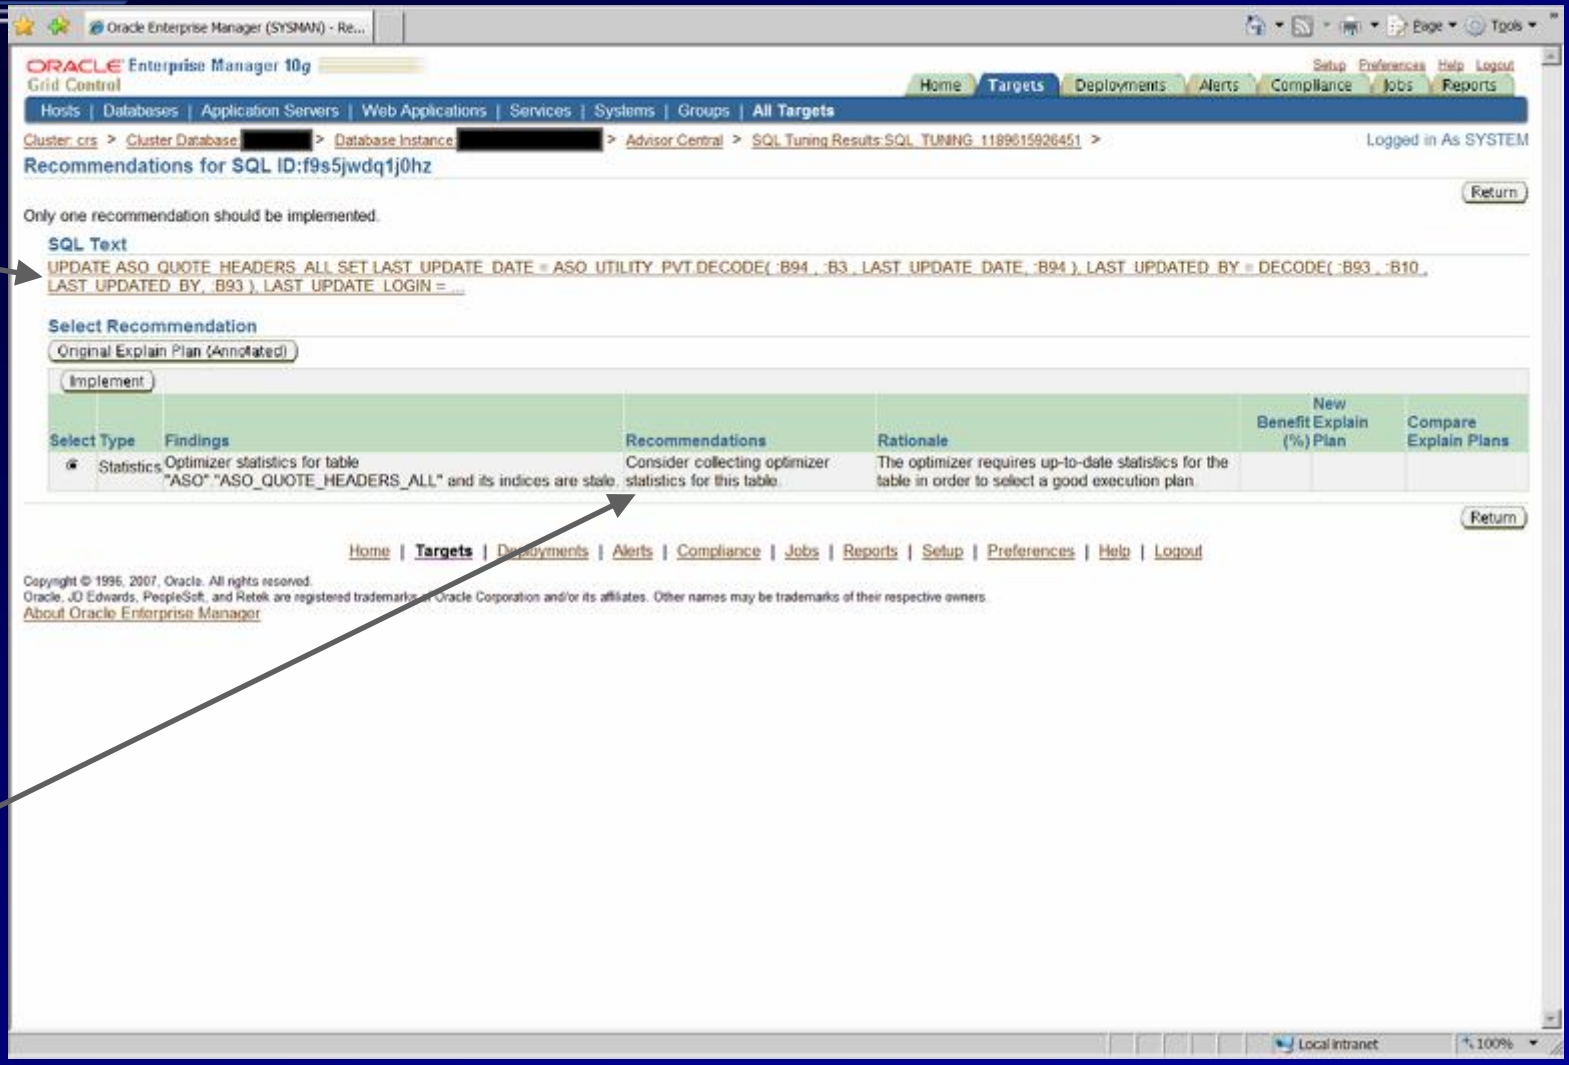

ASH Report – Points to same issues!

ASH Report For (9:12 AM) ... <17 minutes later to verify>

(1 Report Target Specified)

DB Name	DB Id	Instance	Inst num	Release	RAC	Host
DBRJN	277251124	ORA10	2	10.2.0.3.0	YES	linux2

CPUs	SGA Size	Buffer Cache	Shared Pool	ASH Buffer
8	6,000M (100%)	1,536M (25.6%)	1,025M (17.1%)	29.0M (0.5%)

Top User Events

Event	Event Class	% Activity	Avg Active Sessions
enq: TX - row lock contention	Application	99.33	146.11

Top SQL Statements

SQL ID	Planhash	% Activity	Event	% Event
SQL Text				

Go to Hang Analysis & Verify the Pain!
<stay ahead of the problem>

All these Red & Yellow Colors are NOT a good sign!

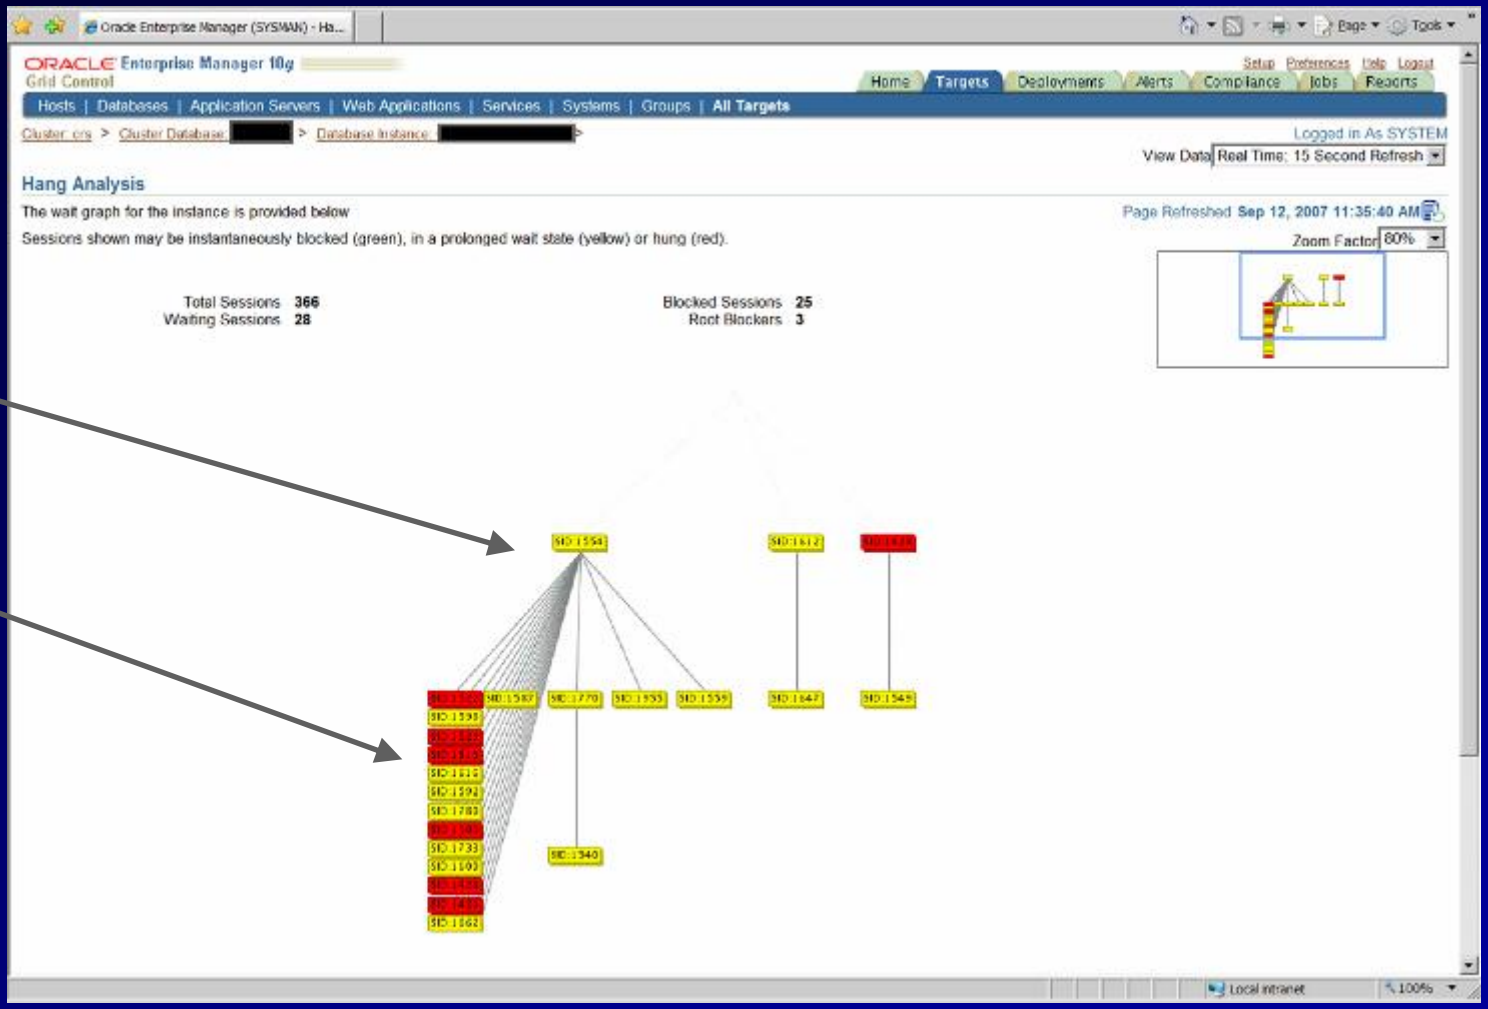

Tuning Multiple Nodes

5 Nodes All Up!

2 Nodes Down 1 Starting up!

Are all the Nodes up in the Cluster?

Tuning the RAC Cluster Interconnect

- Guidelines for GES Statistics:
  - All times should be < 15ms
  - Ratio of global lock gets vs global lock releases should be near 1.0
- High values could indicate possible network or memory problems
- Could also be caused by application locking issues
- May need to review the enqueue section of STATSPACK report for further analysis.

Complete Presentation by Oracle's Rich Niemiec's at:
http://www.oracleracsig.org

Tuning the RAC Cluster Interconnect
Using AWR Reports (FYI Only)

Global Cache Load Profile

	Per Second	Per Transaction
Global Cache blocks received:	0.38	0.05
Global Cache blocks served:	0.26	0.04
GCS/GES messages received:	766.83	106.40
GCS/GES messages sent:	1,278.25	177.36
DBWR Fusion writes:	0.01	0.00
Estd Interconnect traffic (KB)	404.57	

Global Cache Efficiency Percentages (Target local+remote 100%)

Buffer access - local cache %:	100.00
Buffer access - remote cache %:	0.00
Buffer access - disk %:	0.00

Global Cache and Enqueue Services - Workload Characteristics

Avg global enqueue get time (ms):	1.9
Avg global cache cr block receive time (ms):	1.8
Avg global cache current block receive time (ms):	1.9
Avg global cache cr block build time (ms):	0.0
Avg global cache cr block send time (ms):	0.2
Global cache log flushes for cr blocks served %:	0.0
Avg global cache cr block flush time (ms):	
Avg global cache current block pin time (ms):	0.1

Tuning the RAC Cluster Interconnect
Using AWR Reports (FYI Only)

SQL ordered by Cluster Wait Time

Cluster Wait Time (s)	CWT % of Elapsd Time	Elapsed Time(s)	CPU Time (s)	Executions	SQL Id	SQL Module	SQL Text
75.00	0.38	19,515.02	7,568.17	1,680			
0.85	3.63	23.36	7.56	1			
0.15	13.69	1.09	0.76	1			
0.09	31.73	0.27	0.08	46			
0.06	12.31	0.46	0.30	1			
0.03	0.20	14.19	6.18	1,688			
0.03	0.26	10.71	2.71	1,688			
0.03	83.42	0.03	0.00	1			
0.02	36.30	0.07	0.06	4			
0.02	43.85	0.06	0.01	1			
0.02	1.77	1.12	0.54	10			
0.02	10.23	0.15	0.13	1			
0.02	3.76	0.40	0.02	1			
0.01	12.17	0.11	0.09	1			
0.01	33.26	0.03	0.02	1			
0.01	8.39	0.09	0.09	177			

Done Internet

Future Tuning – 11g FM

Performance Manager : Back in Time!

OVERVIEW - rich1

Window Refresh Help

Automatic Refresh Interval: 0 min 30 sec

BUFFER CACHE HIT%
4.14%
95.86%
Hit % Miss %

LIBRARY CACHE HIT% ...
99.91%
Hit % Miss %

DATA DICT CACHE HIT%
12.41%
87.59%
Hit % Miss %

MEMORY SORT HIT%
1.89%
98.11%
Hit % Miss %

ROLLBACK NOWAIT HIT%
100%
Hit % Miss %

USERS LOGGED ON
11.00
11.00 8.80 6.60 4.40 2.20 0.00
1

ACTIVE USERS
5.00
6.00 4.80 3.60 2.40 1.20 0.00
1

USERS RUNNING
4.00
4.00 3.20 2.40 1.60 0.80 0.00
1

USERS WAITING FOR LOCKS
0.00
1.00 0.80 0.60 0.40 0.20 0.00
1

MEMORY ALLOCATED ...
688320
688320 550656 412992 275328 137664 0
1

FILE I/O RATE ...
44.00 22.00 0.00
Physical Reads/s Physical Writes/s

SYSTEM I/O RATE
590 295 0
Buffer Gets/s Block Changes/s

THROUGHPUT ...
8.83 0.00
User Calls/s Transactions/s

For Help, press F1

SQL Tuning Advisors & SQL Plan Management (SPM)

SQL Plan Management

- SQL Plan Management is a mechanism that records/evaluates execution plan of SQL statements (good & bad) over time and builds SQL Plan baselines (replaces stored outlines) of existing plans known to be efficient.
- Events that cause the need for SQL Plan baselines:
  - New version of Oracle (New optimizer version – Use capture replay to test effect)
  - Changes to optimizer statistics or data changes
  - Schema, application or metadata changes (use SQL Advisor to get suggestions)
  - System settings changes (Use SQL Replay to find what works)
  - SQL Profile (statistics – data skews & correlated columns) creation
- Stored outlines are deprecated (discouraged) in Oracle Database 11g. Oracle highly recommends migrating existing stored outlines to SQL plan baselines. A **SQL Profile contains additional STATISTICS** for this SQL statement for the query optimizer to generate a better execution plan. **An outline/baseline contains HINTS** for this SQL statement for query optimizer to generate a

SQL Plan Management

- **SQL Profile stores STATISTICS** for a SQL statement for the query optimizer to generate a better execution plan.
- **A Stored Outline/SQL Plan Baseline contains HINTS** for this SQL statement for query optimizer to generate a better execution plan.
- A SQL Plan Baseline should evolve with changes in the system to analyze good/bad plans over time.
- View these in DBA_PLAN_BASELINES
- You can also export a SQL Tuning Set and import it to new system. Capture baselines for Tuning Set with DBMS_SPM (see later slide on entire syntax). Can also use a pack/unpack function to pack/unpack all plans in a system for transporting

SQL Plan Management
Create a SQL Tuning Set

Tuning Issue

Create a Tuning Set from Top 10 SQL

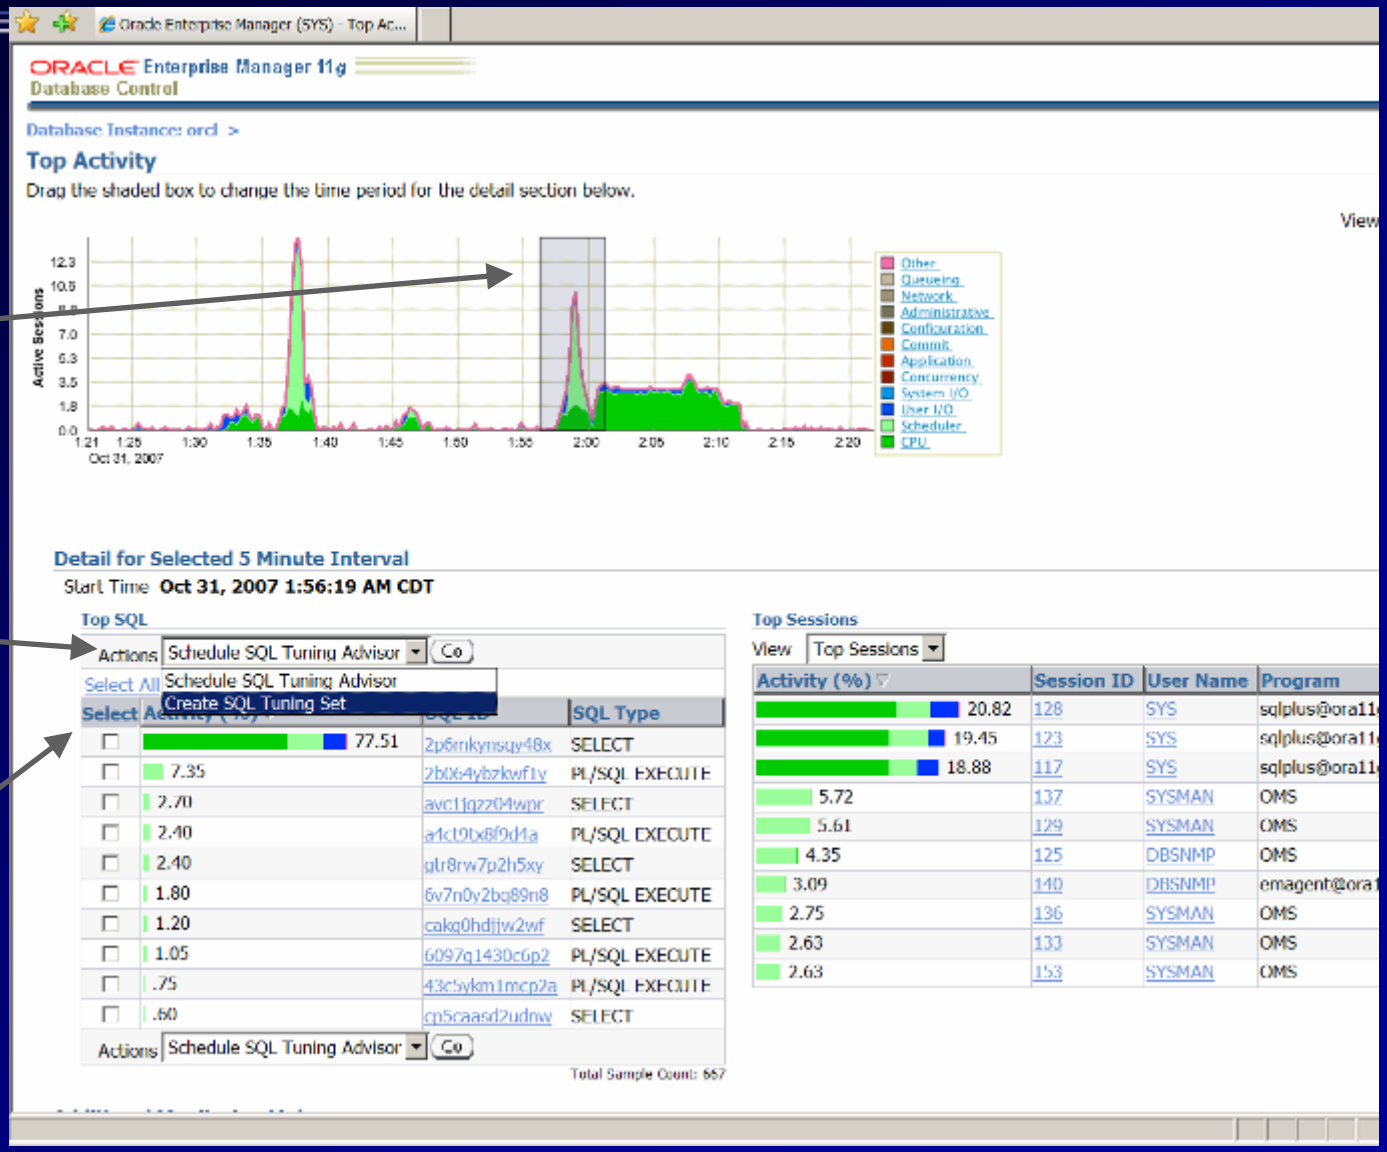

Top Activity – 11gR2 (same look)

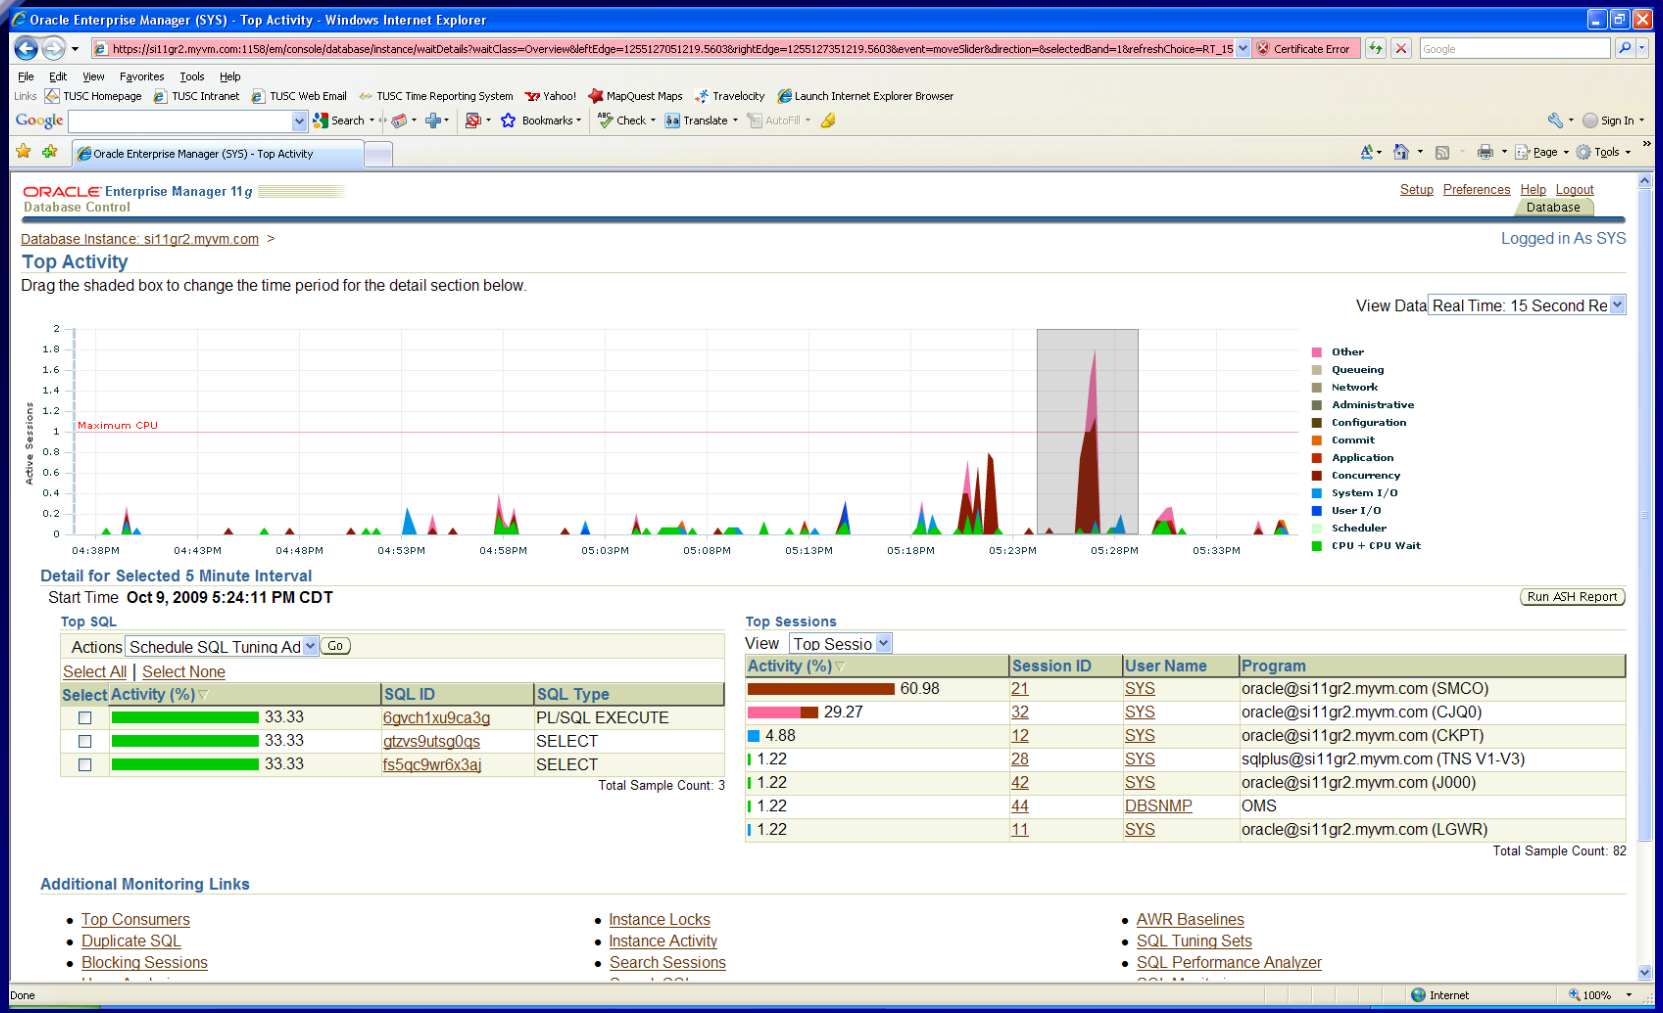

SQL Plan Management
Create a SQL Tuning Set

Tuning Set Name

Queries

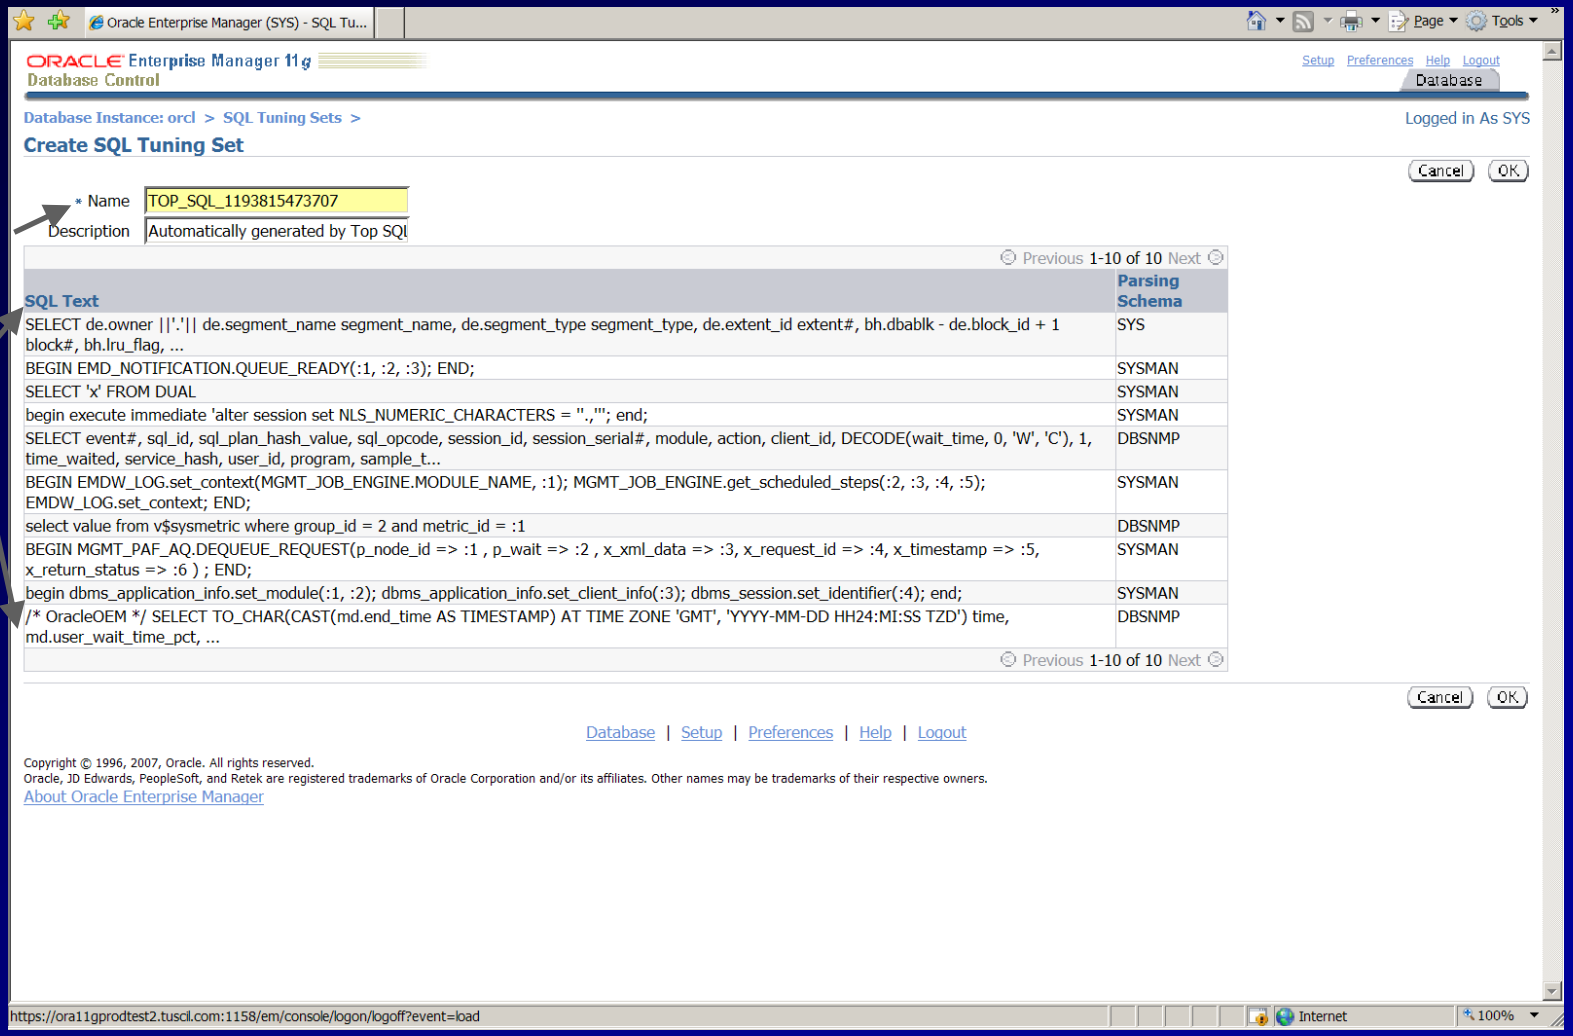

SQL Plan Management
Viewing a SQL Tuning Set

Tuning Set Name

Queries & Stats

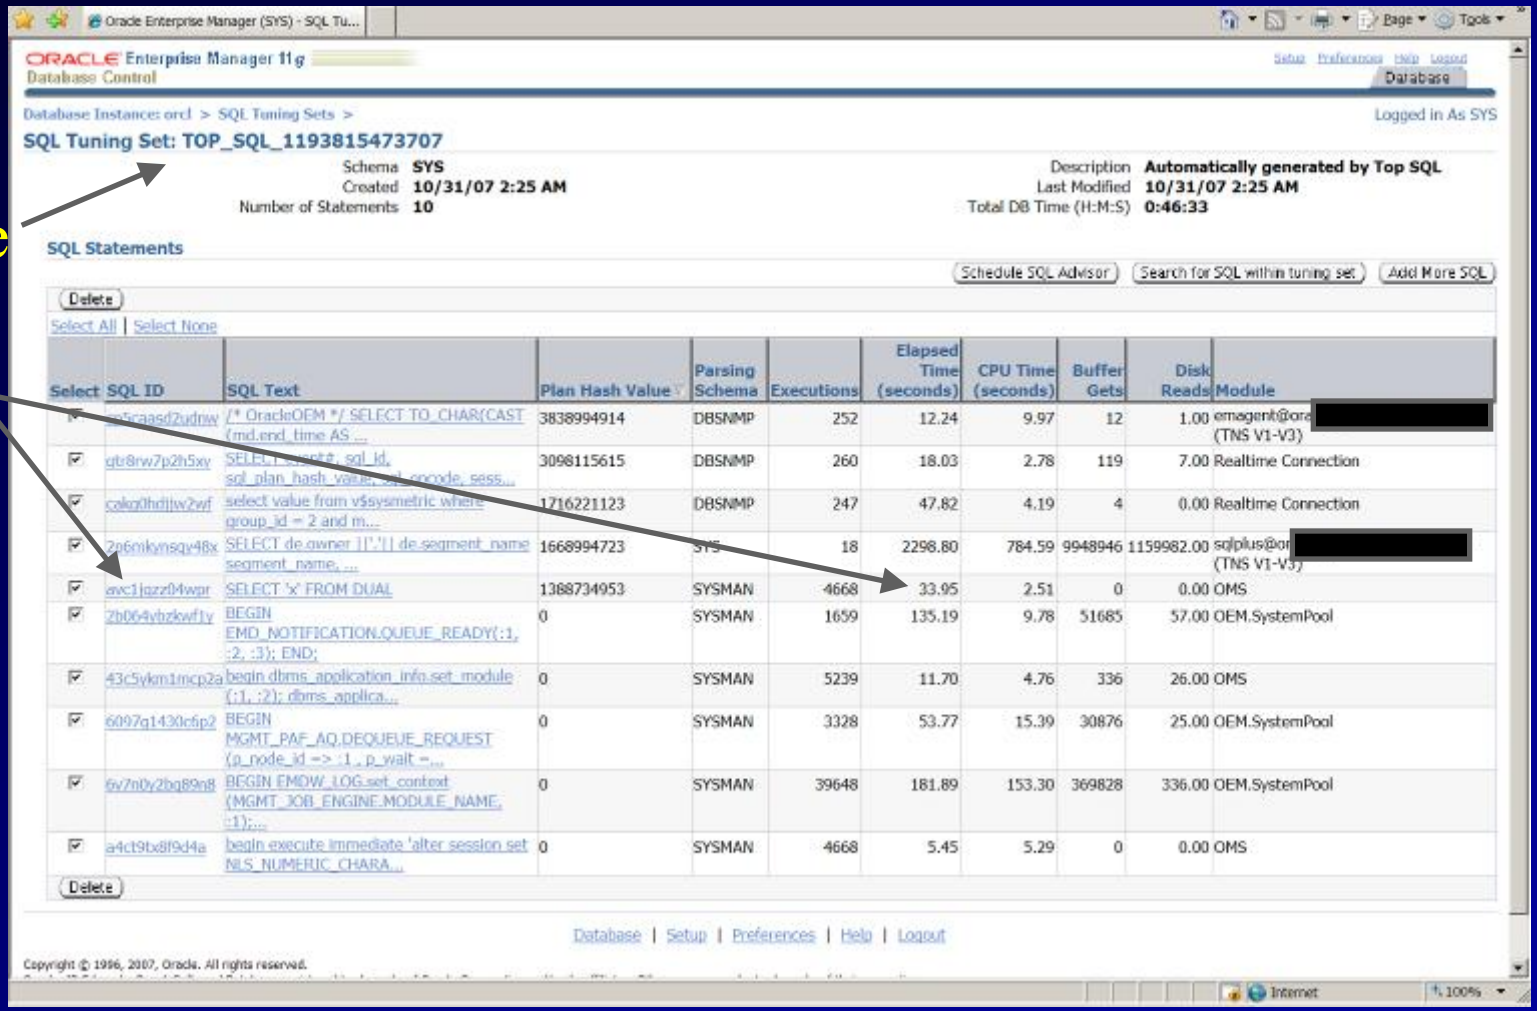

SQL Plan Management
Create a SQL Tuning Set

Run the Tuning Advisor on this SQL Tuning Set (STS)

Run it NOW

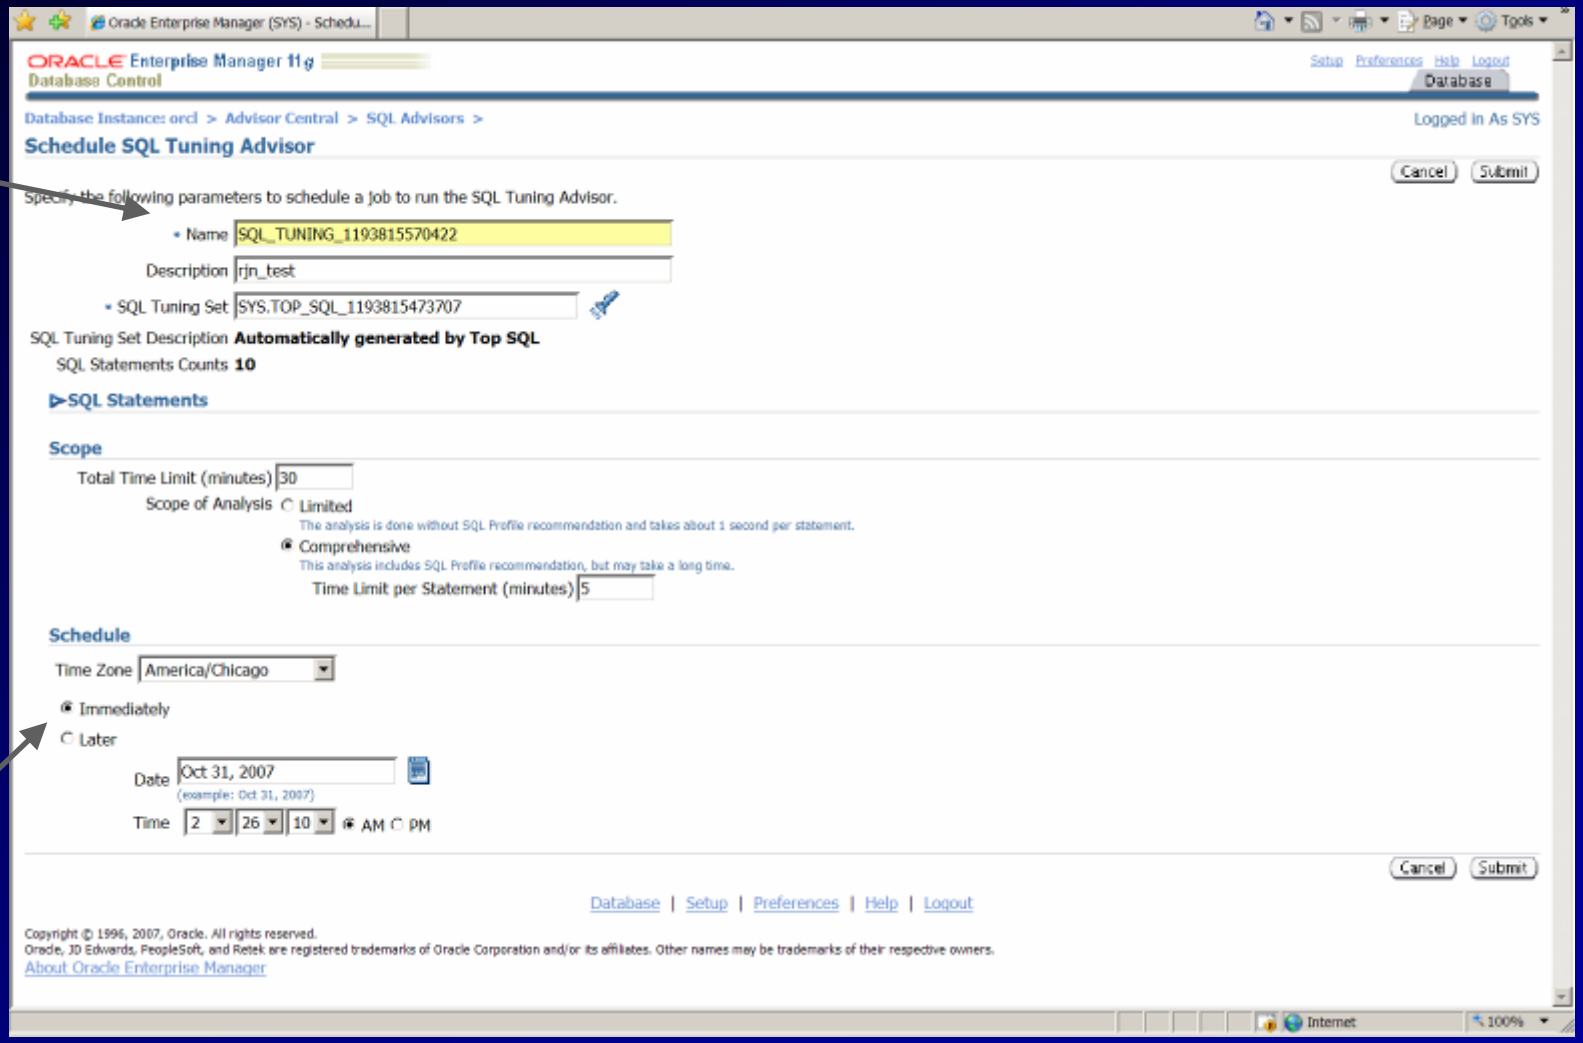

SQL Plan Management
Create a SQL Tuning Set

Results

Select
One query
And click
View

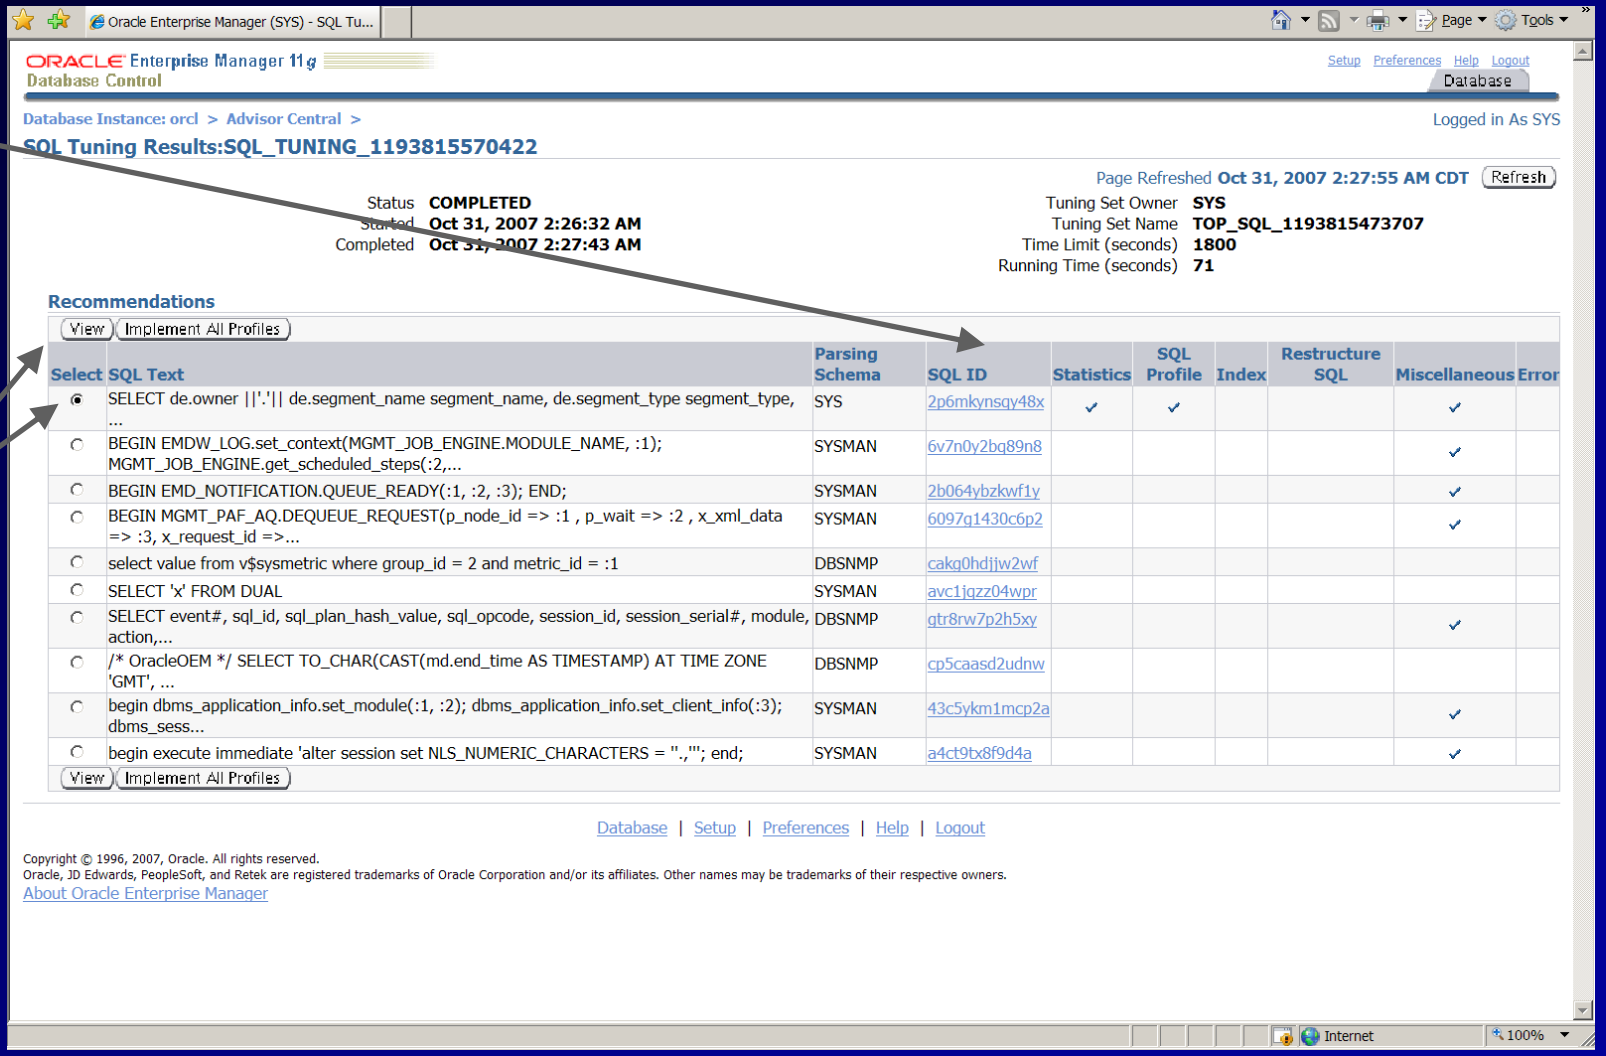

SQL Plan Management
Click on any SQL ID

SQL Text

Waits & Statistics

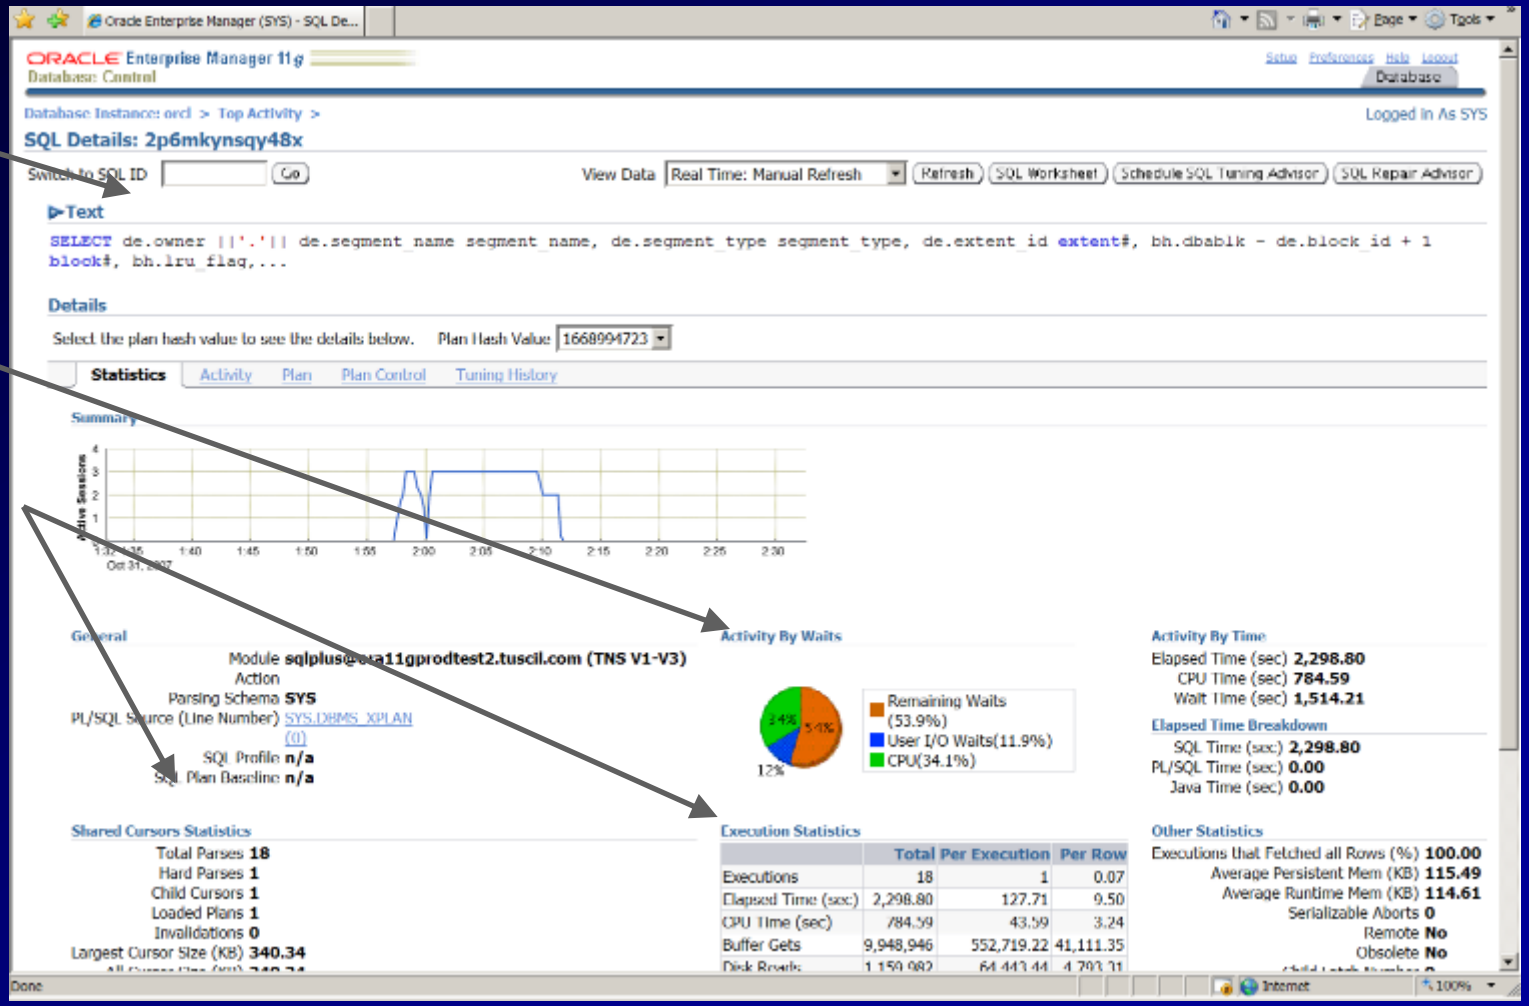

SQL Plan Management
Create a SQL Tuning Set

SQL Profile Will Help 99%

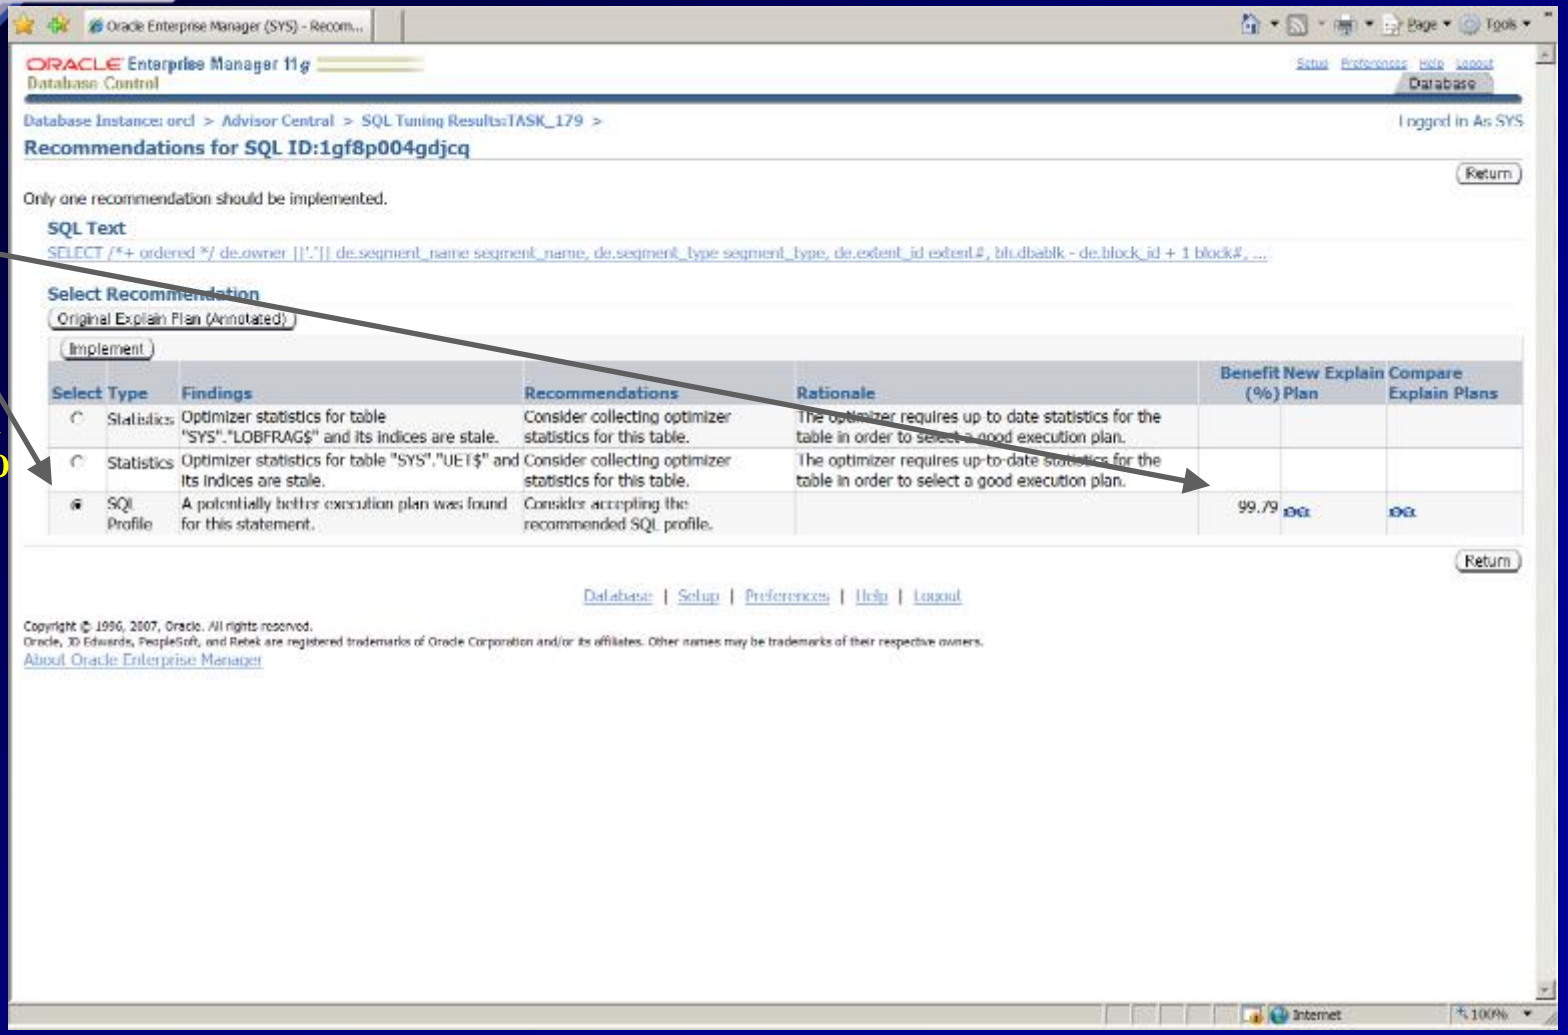

SQL Plan Management
Create a SQL Tuning Set

Compare Before & After

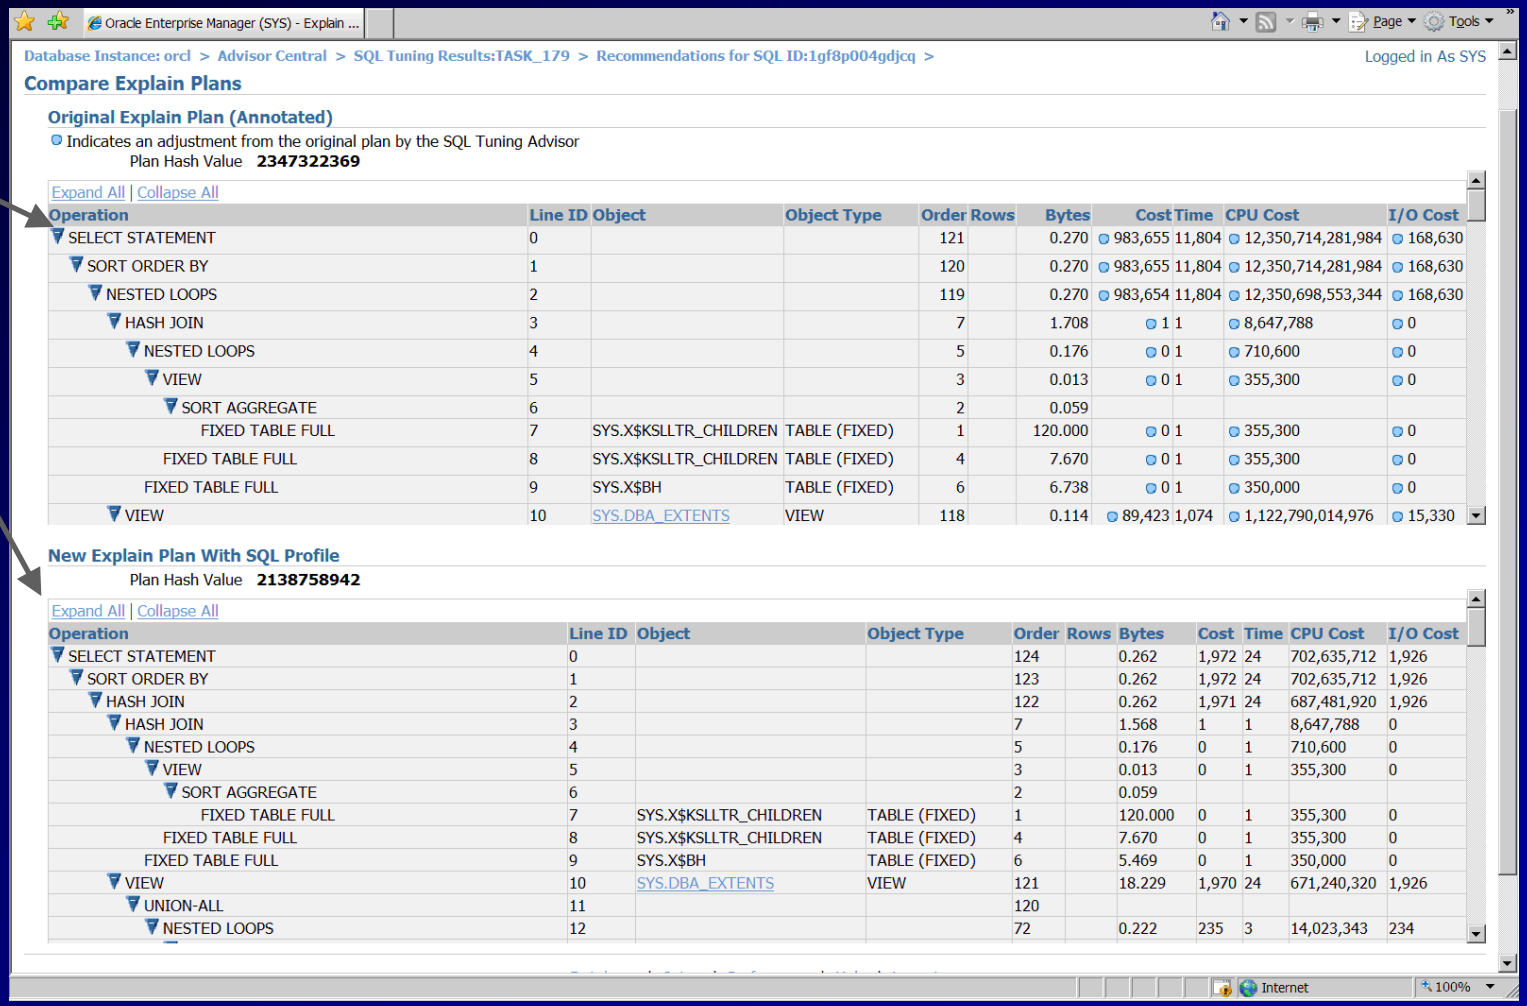

SQL Plan Control
SQL Profiles stored in the system

SQL Profiles

SQL Plan Baselines

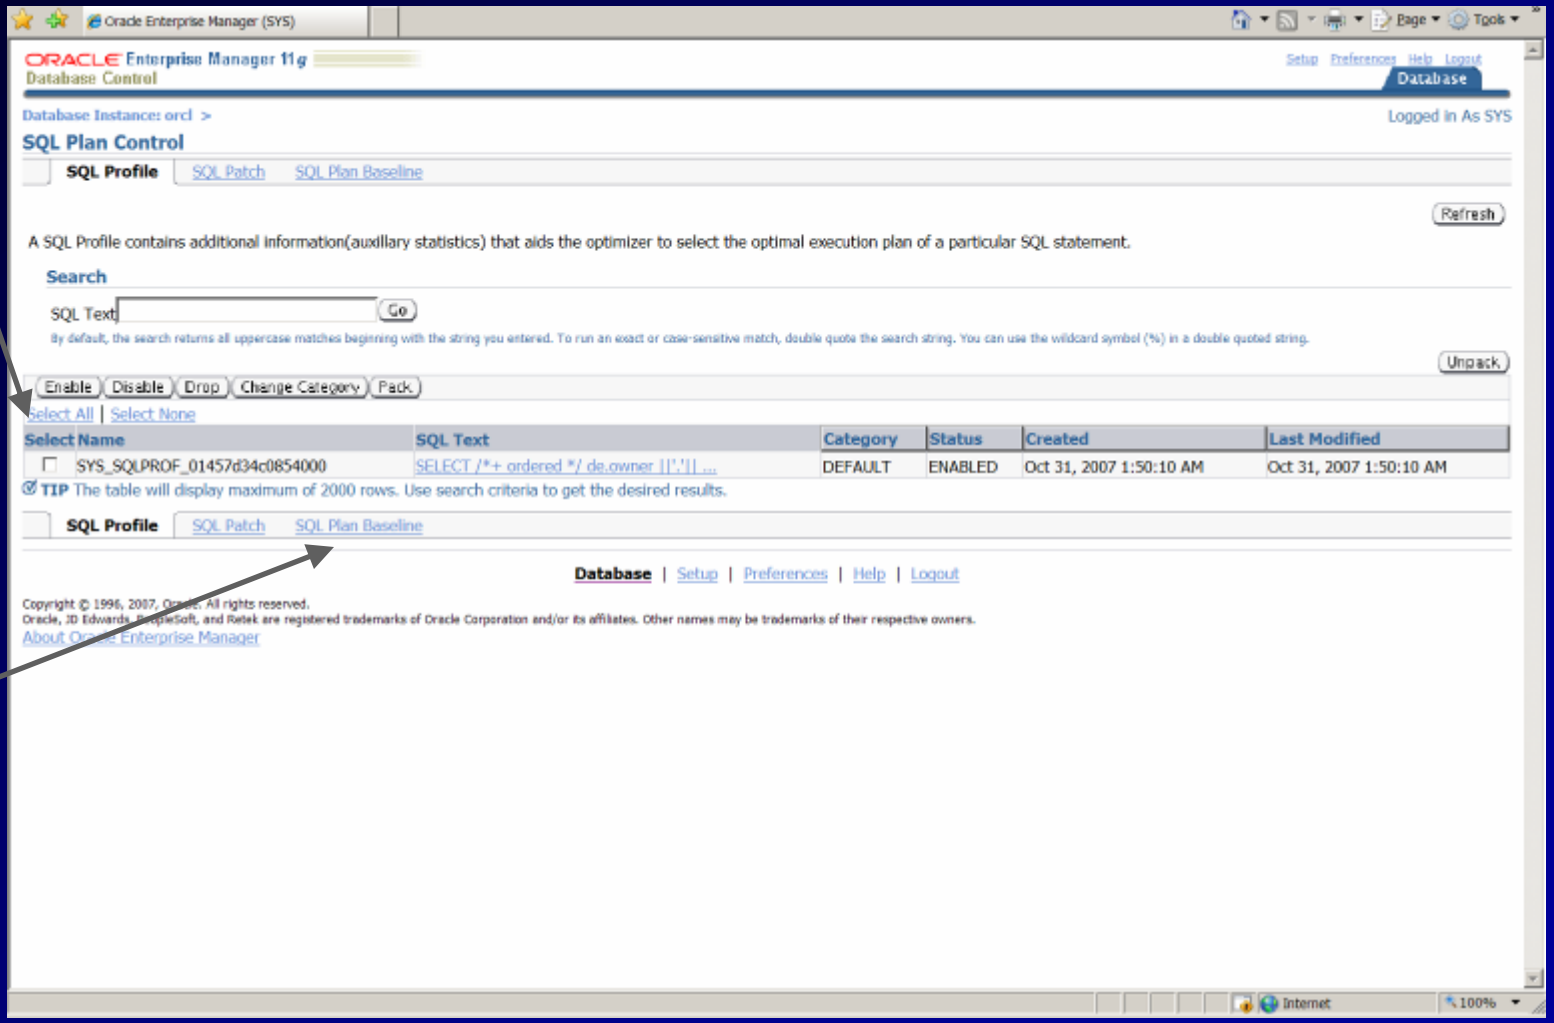

SQL Performance Analyzer

SQL Performance Analyzer

- Measure and report on performance before and after a change! DBMS_SQLTUNE package.

Great for:

- Database Upgrades
- Application Upgrades
- Hardware Changes
- Database or Schema Changes
- Best for SQL Tuning – Especially Batches

SQL Performance Analyzer

Easy to run – SQL Focus (Test SGA settings):

- Capture SQL
- Transport SQL
- Create a Replay Task
- Set up the environment to Test
- Make any changes to Test (such as SGA/Optimizer)
- Compare before and after performance
- Tune the problems!

SQL Performance Analyzer

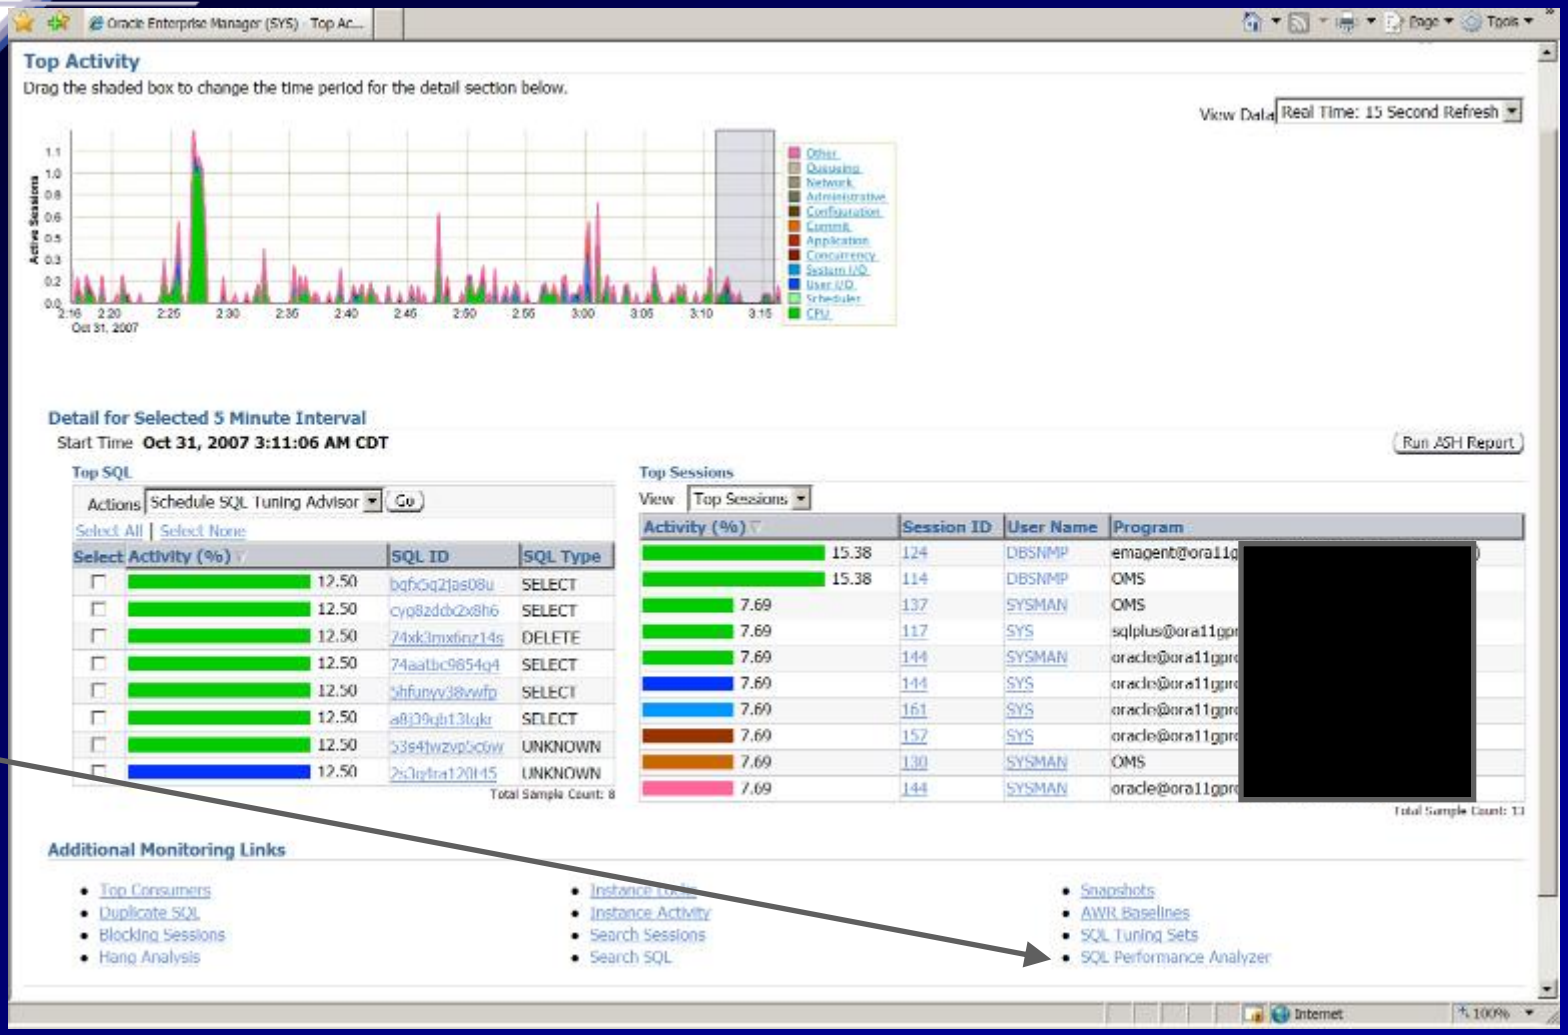

Click Here

SQL Performance Analyzer Guided Workflow

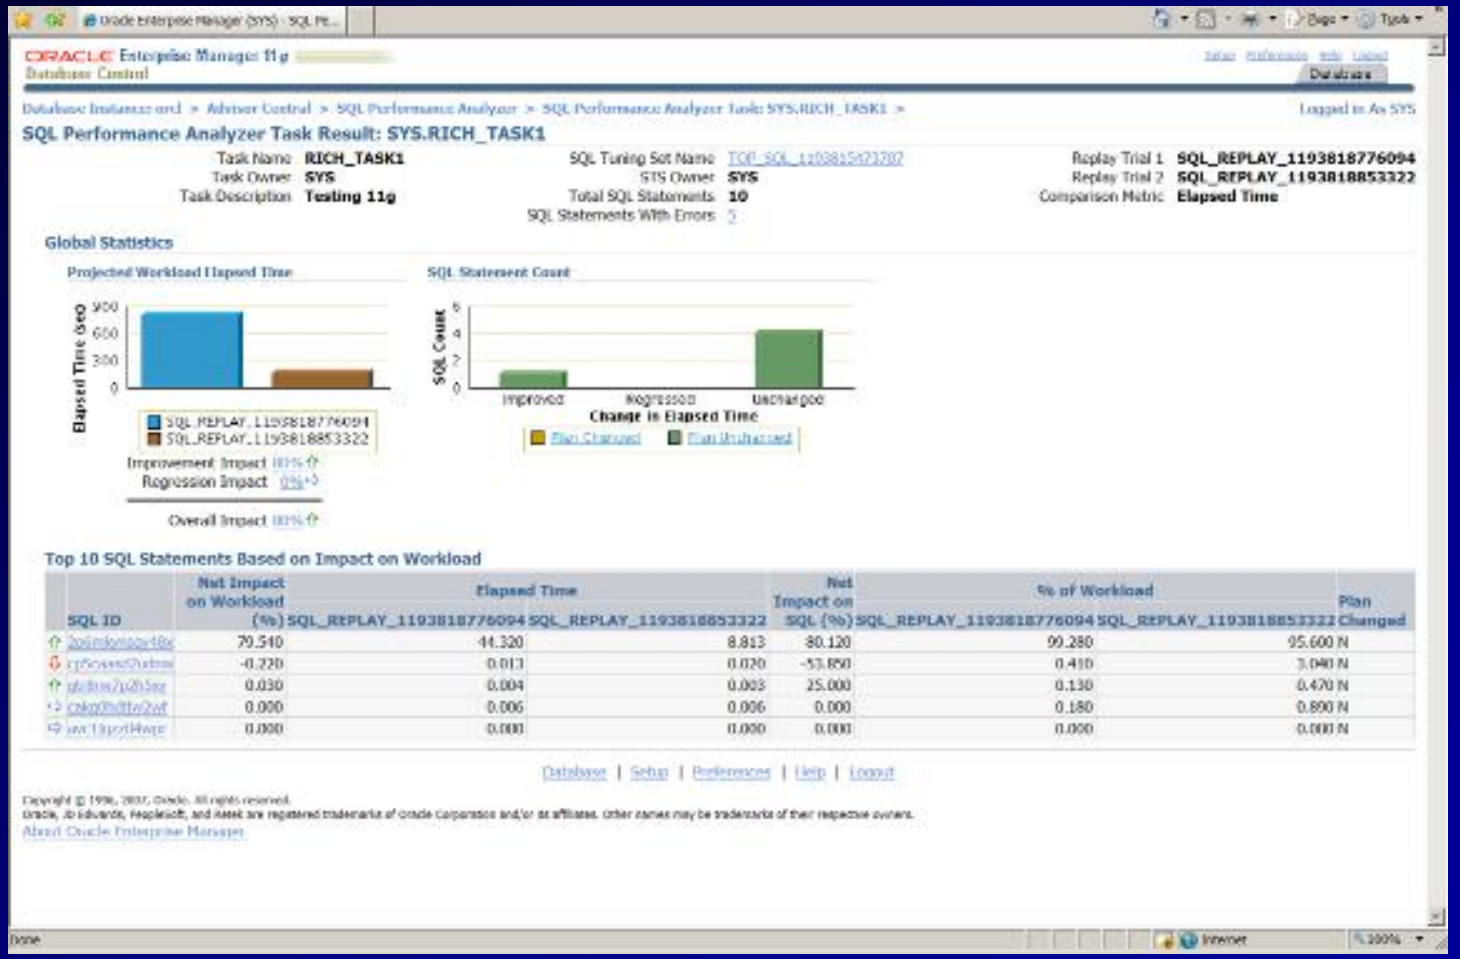

SQL Performance Analyzer Optimizer Upgrade (10g vs. 11g)

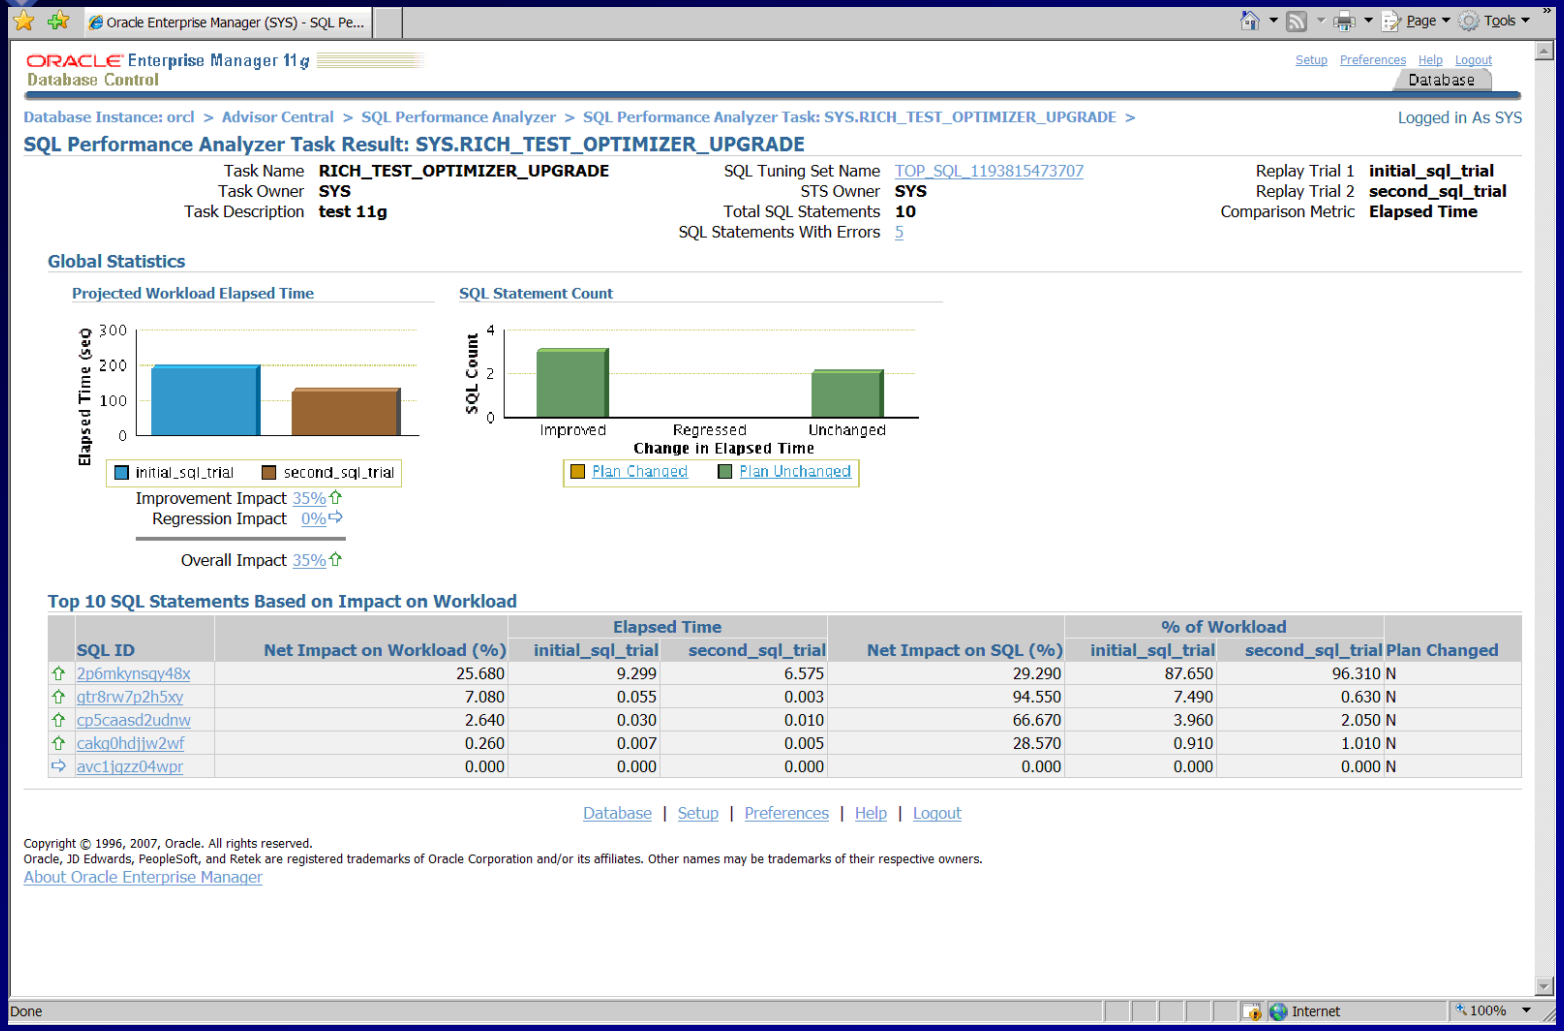

SQL Performance Analyzer 11gR2 - Options

Upgrade Options

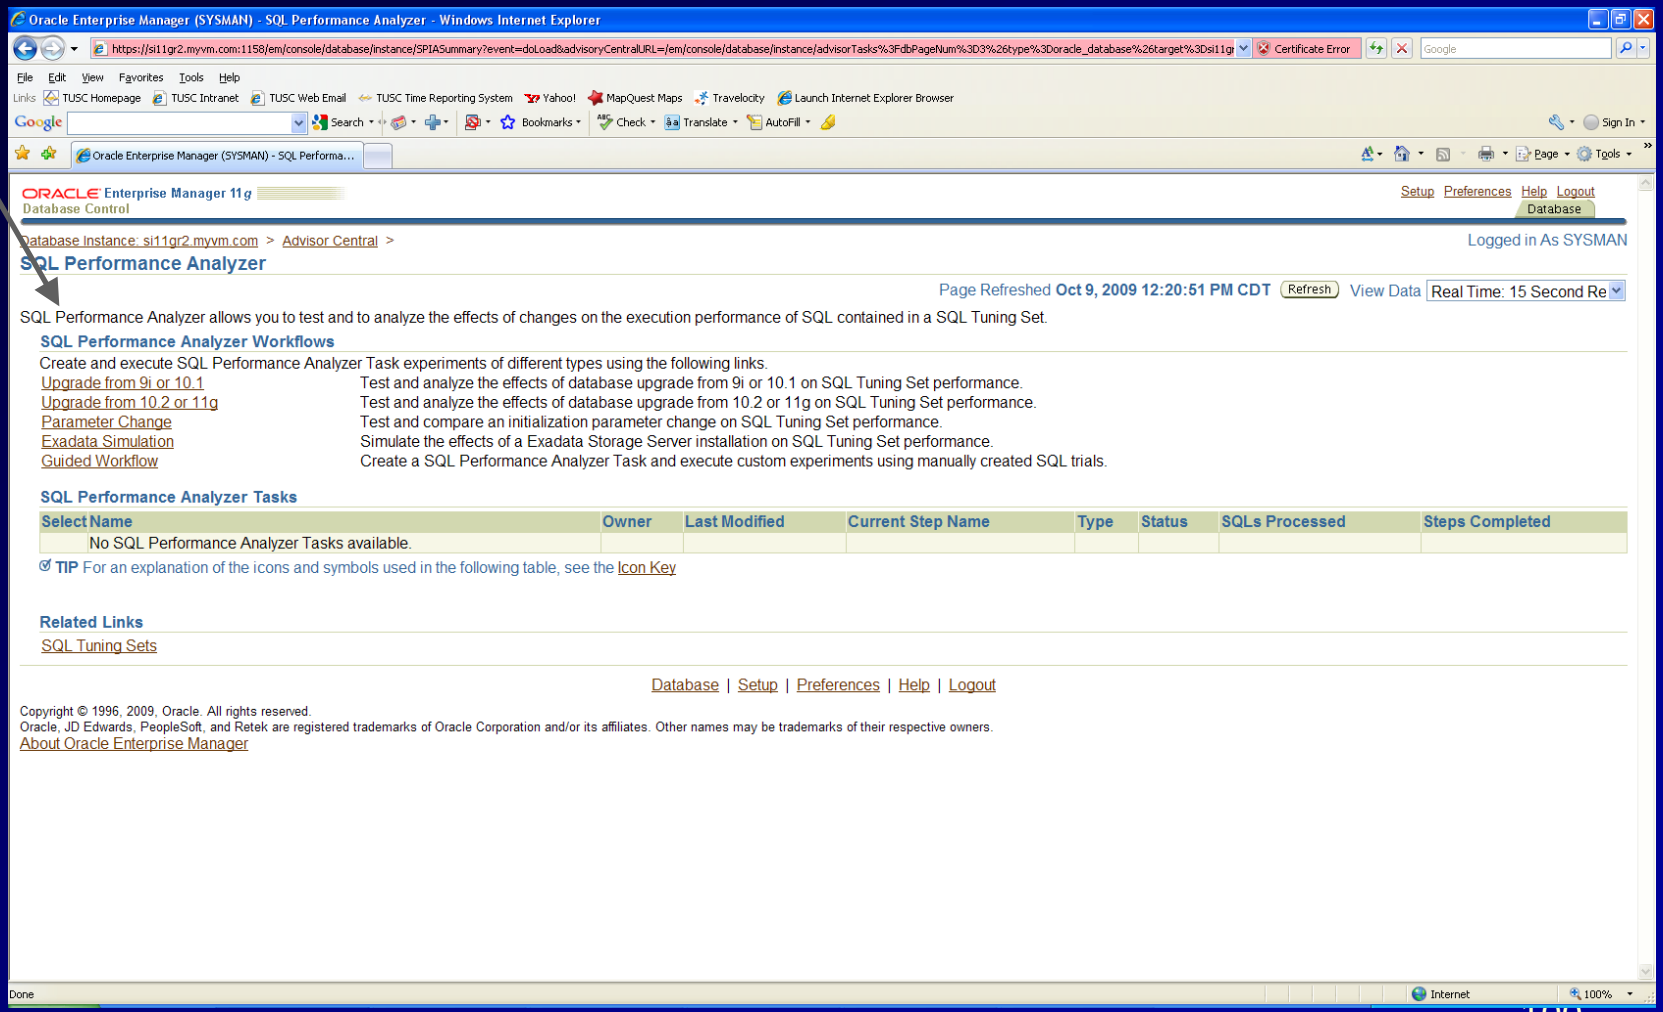

SQL Performance Analyzer 11gR2 – Exadata Simulation

Simple Job so no benefit

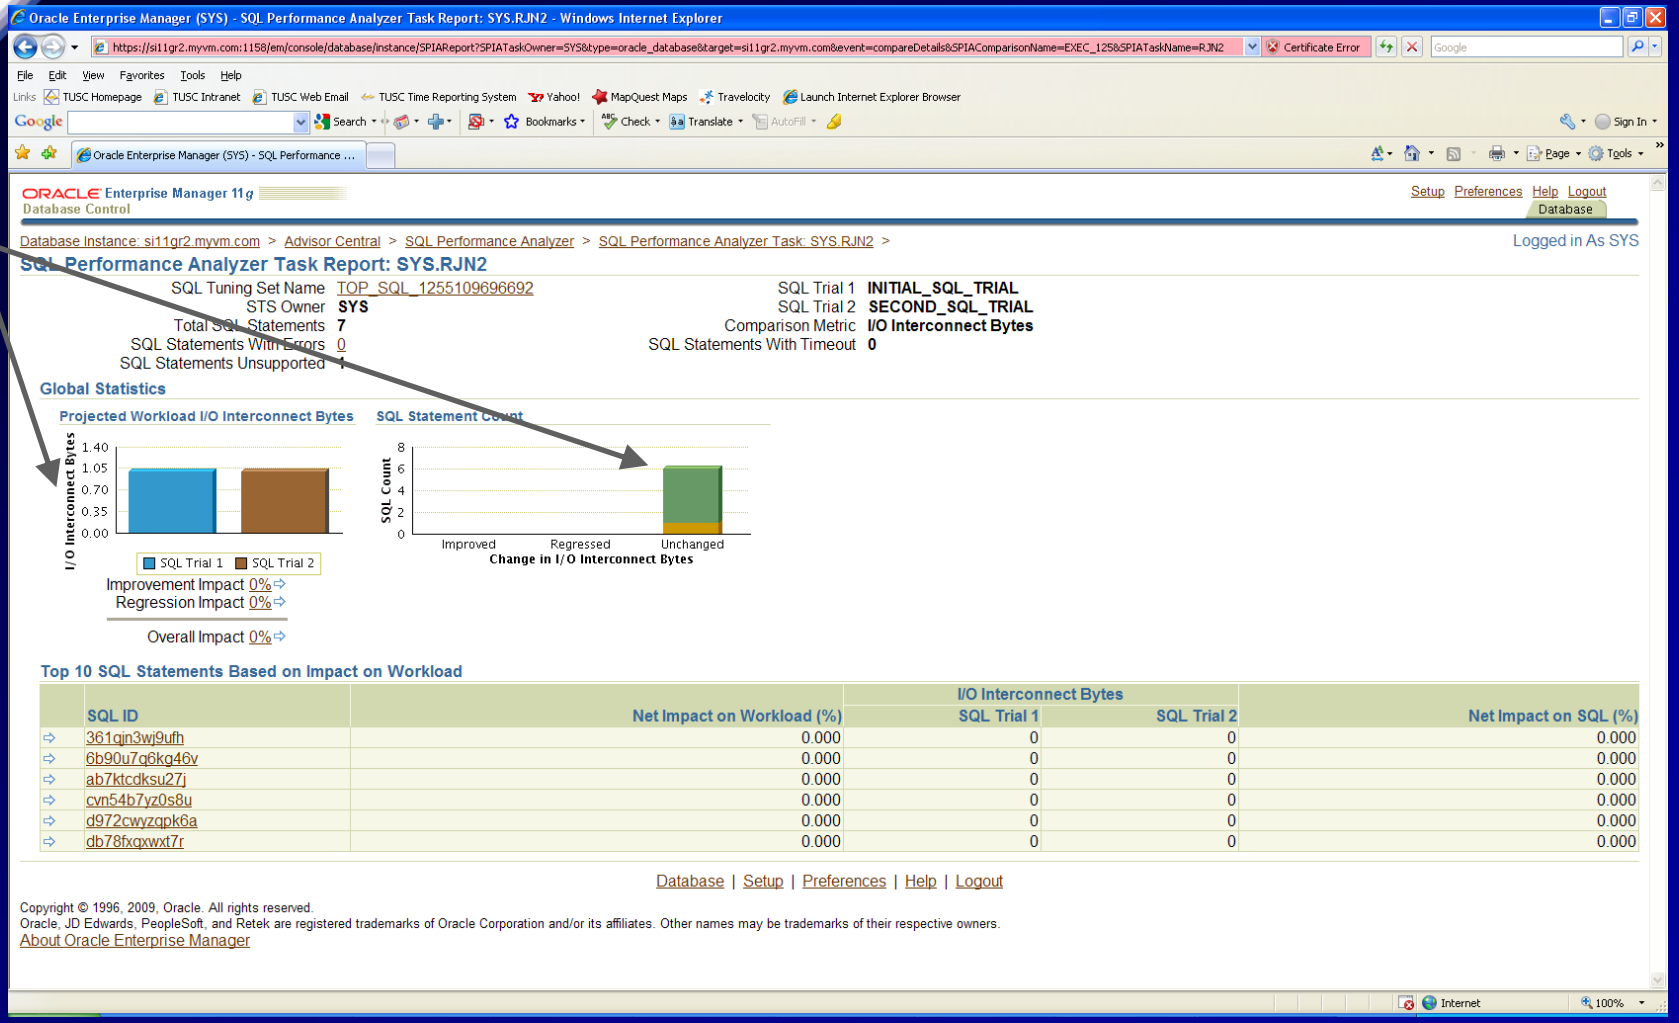

SQL Advisors

Tuning Advisors

Repair Advisor (next)

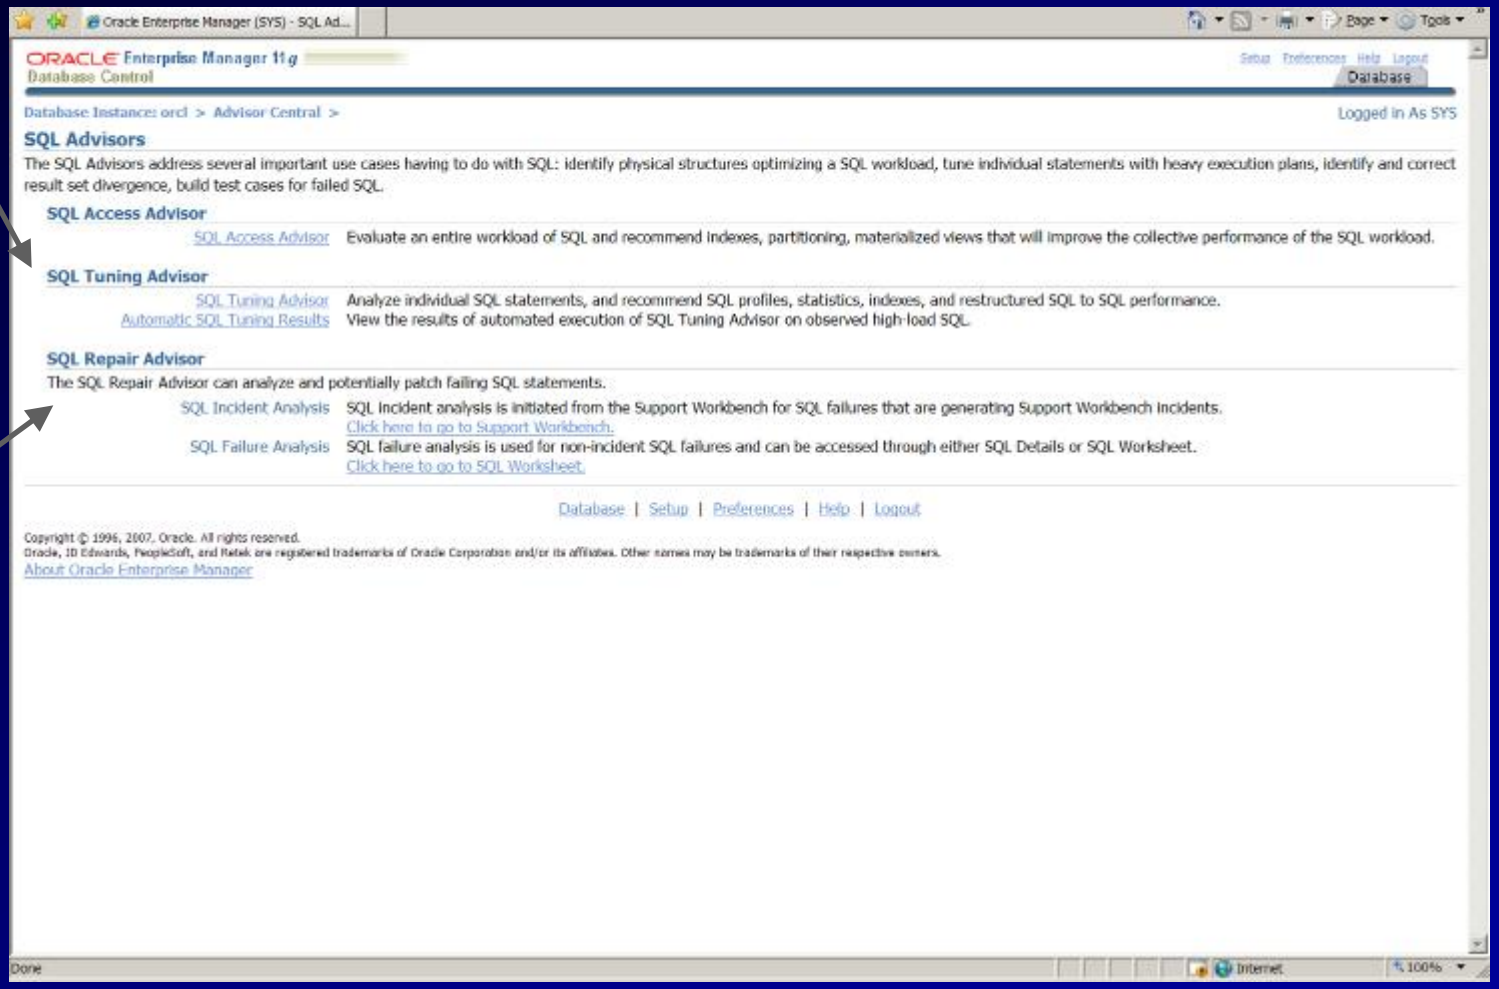

SQL Advisors
11gR2 – (same)

Tuning Advisors

Repair Advisor (next)

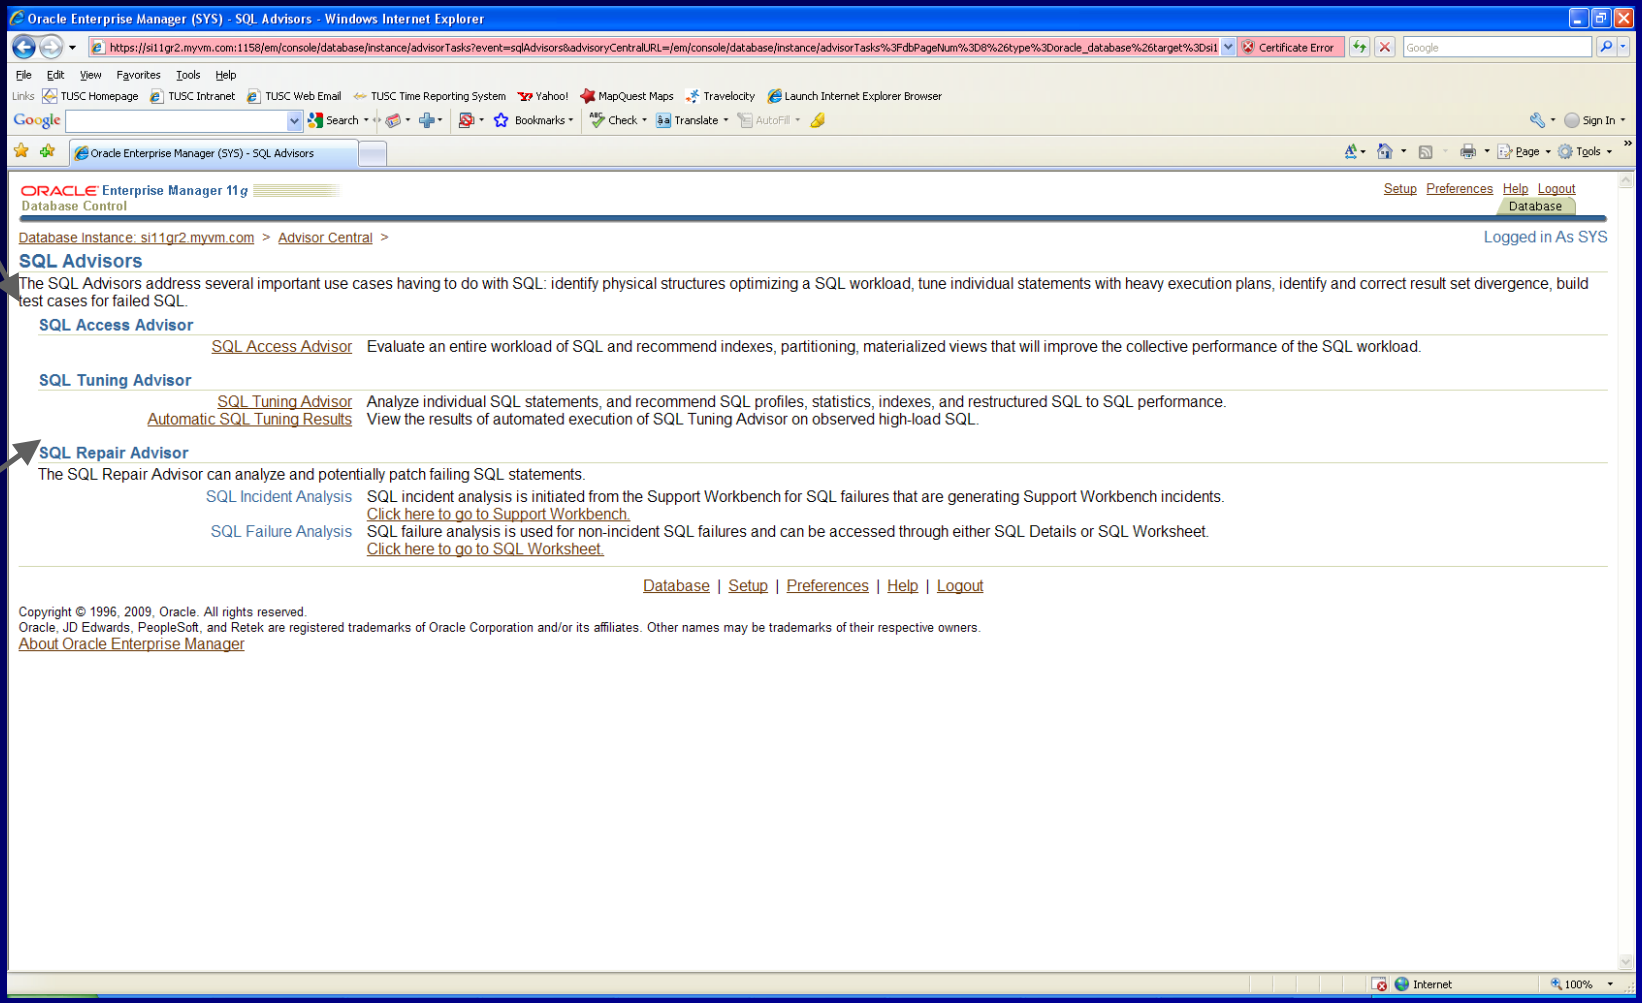

SQL Access Advisor
NEW Partition Advisor

SQL Access Advisor & NEW Partition Advisor

- The SQL Advisor now combines the functionality of the SQL Tuning Advisor, SQL Access Advisor and the new Partition Advisor.
  - Recommends Partitioning Needs
  - Utilize a previous SQL Tuning Set
  - Take SQL straight from what's currently in the CACHE.
  - Create a hypothetical workload
  - SQL Access Advisor checks Indexes, Partitions or Materialized Views (**schema related issues**)

SQL Access Advisor & NEW Partition Advisor

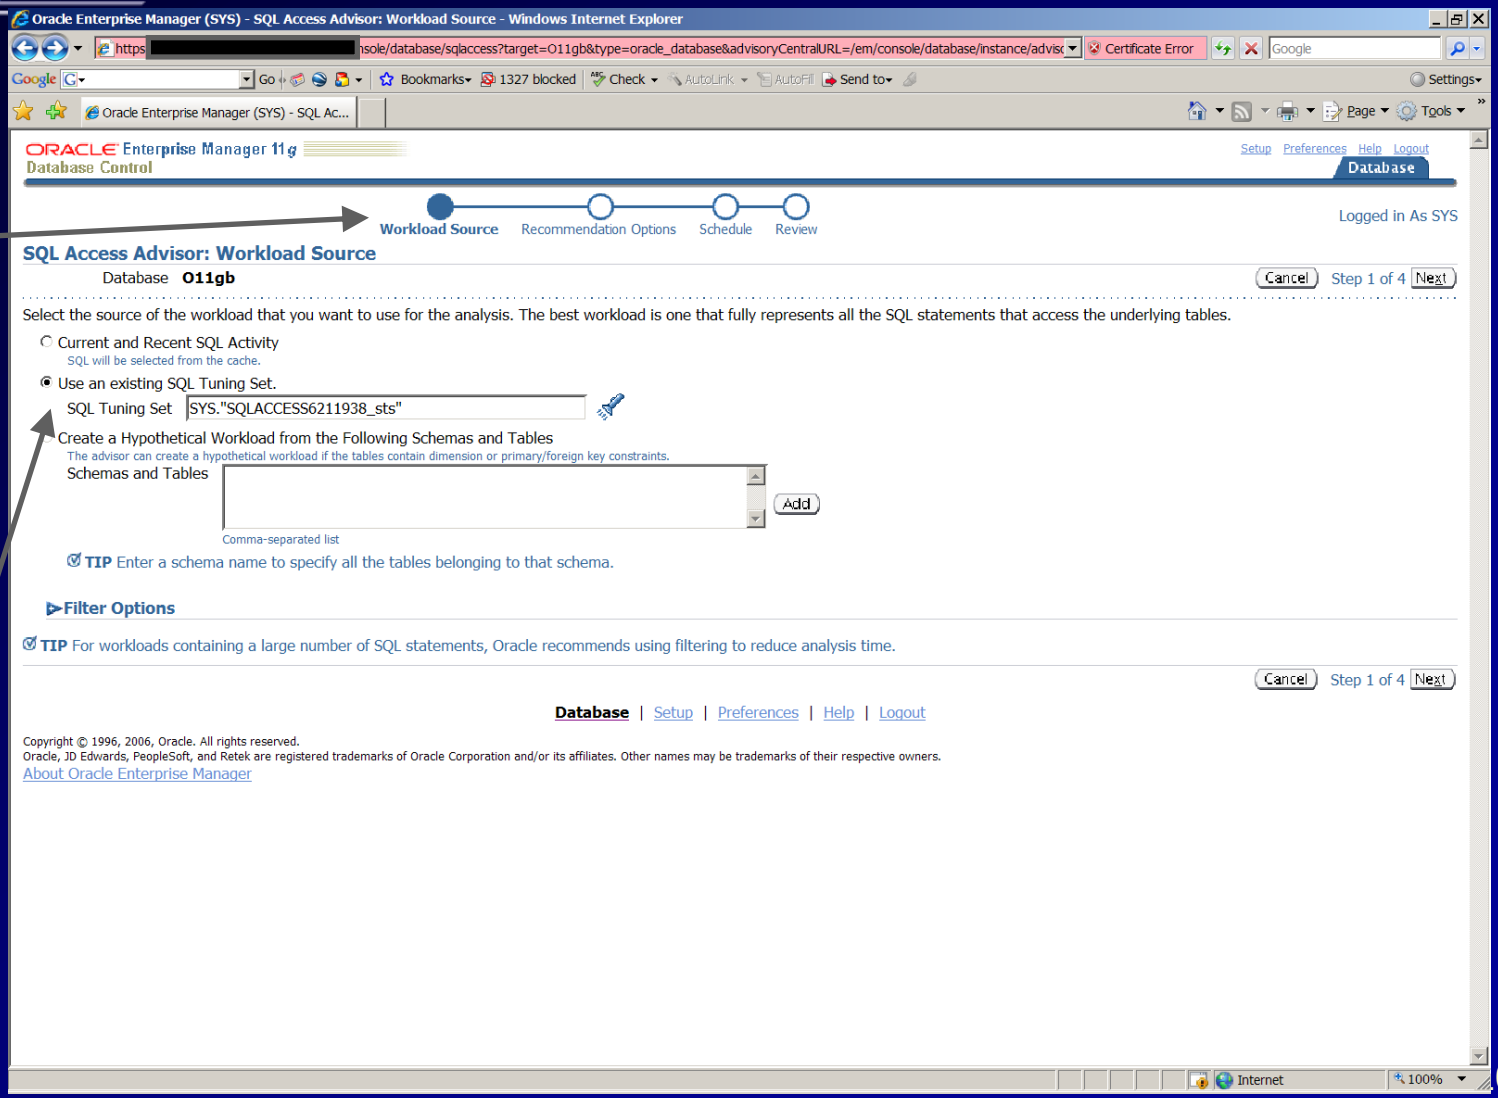

Step One

Use a SQL Tuning Set

SQL Access Advisor & NEW Partition Advisor

Look at Partitions

Quick Solution

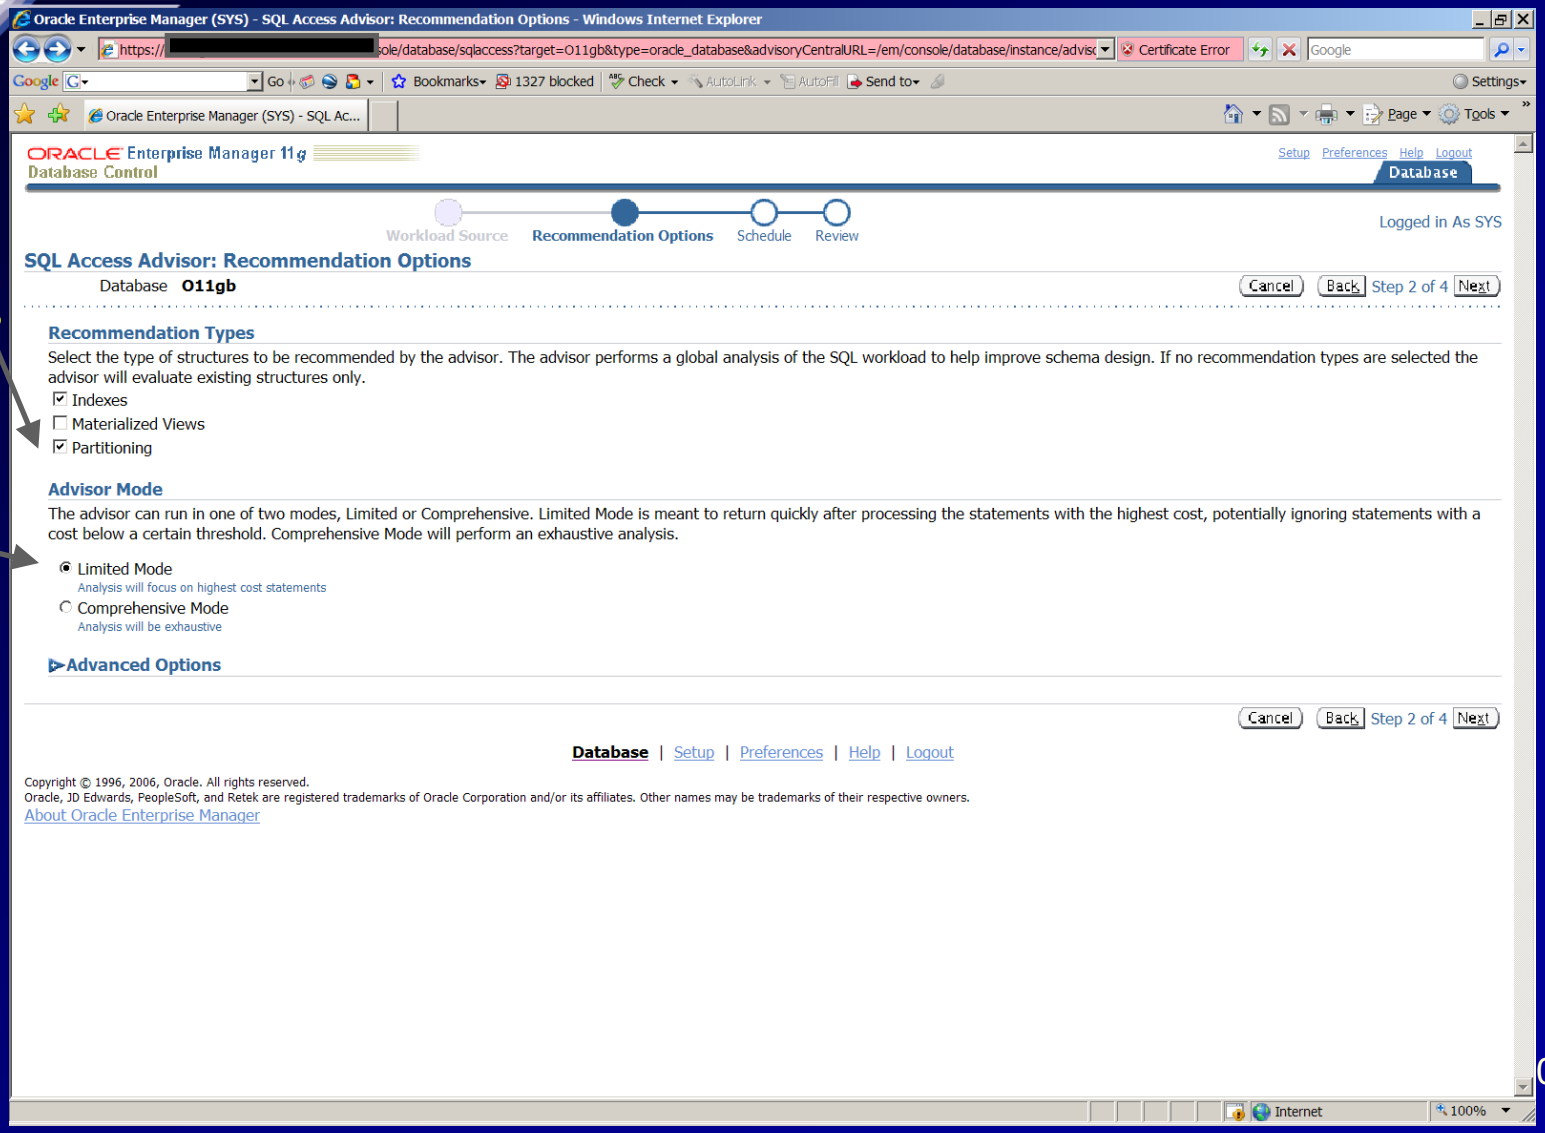

SQL Access Advisor & NEW Partition Advisor

Improve I/O

Improve Execution Time

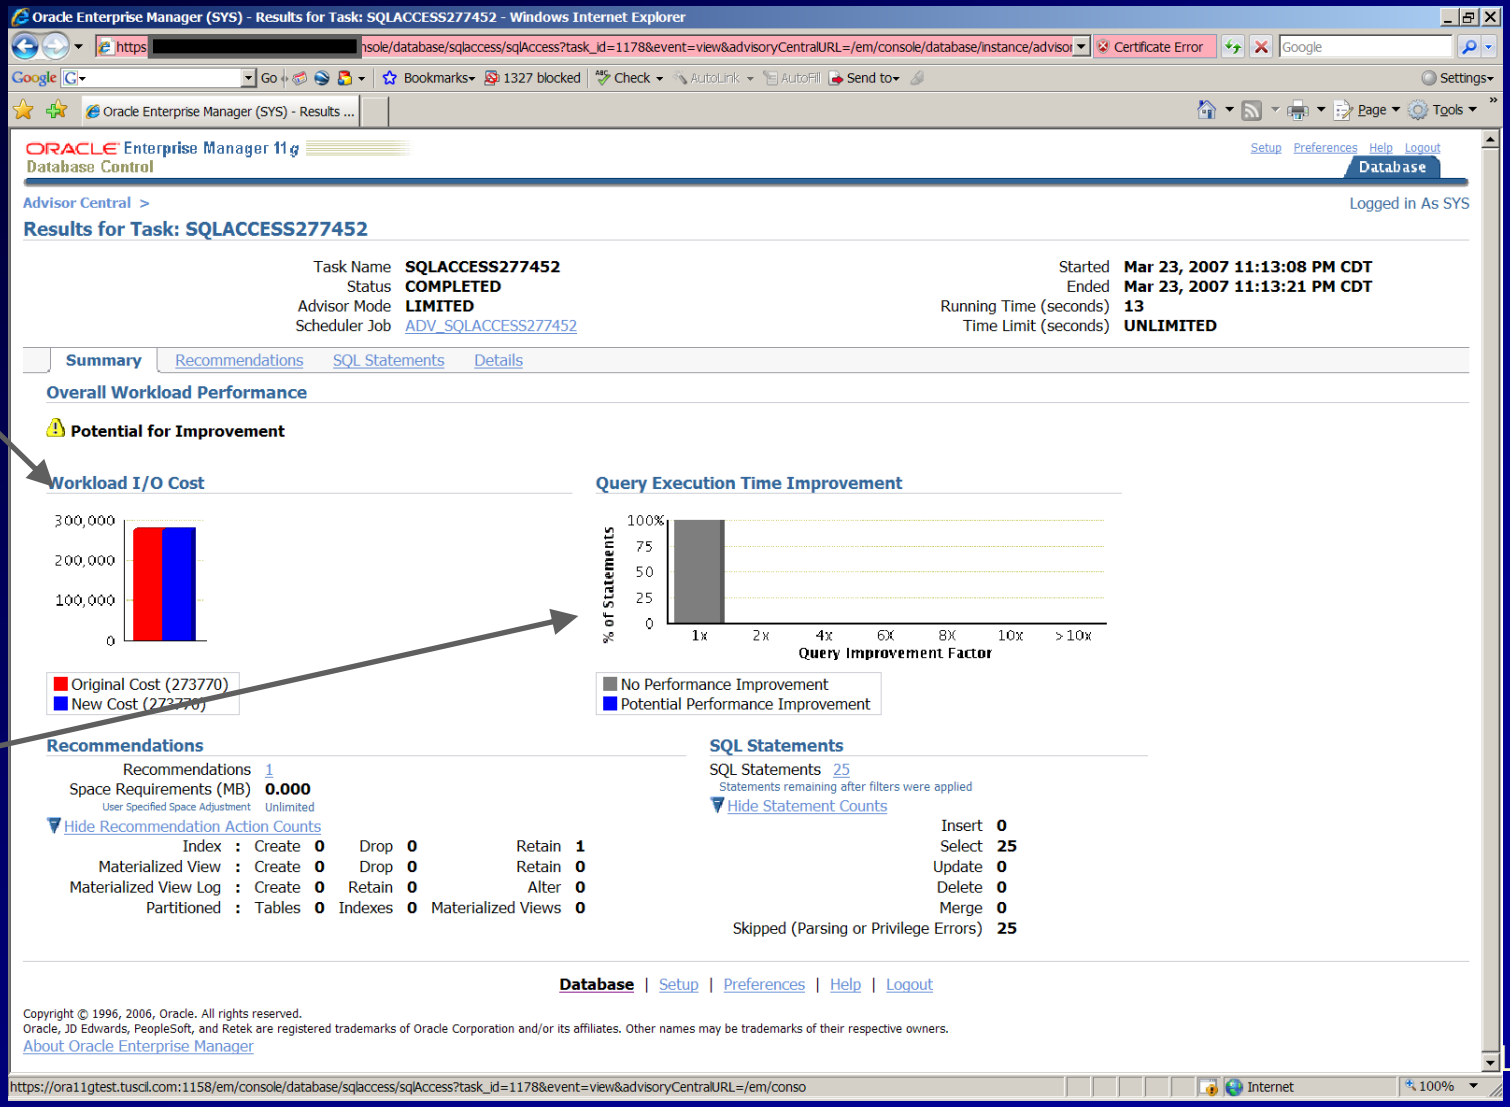

The SQL Repair Advisor

Repair the Problem "*on the fly*"

The Business of IT is serving information…
Not giving users ORA-600 errors…

SQL Repair Advisor

- Used to Repair Problem SQL – Oracle Errors
- Reloads and recompiles SQL statements to gather diagnostics information to fix.
- Uses the diagnostic information to repair the problem SQL statement (DBMS_SQLDIAG)
- Will fix error **going through compilation, execution and trying different routes** (could be a slower route for now) to come up with a temporary SQL Patch without error until fixed.

SQL Repair Advisor – Go straight from Alerts

Go to the Database Instance

Click Alert (ORA-600) message text to see details

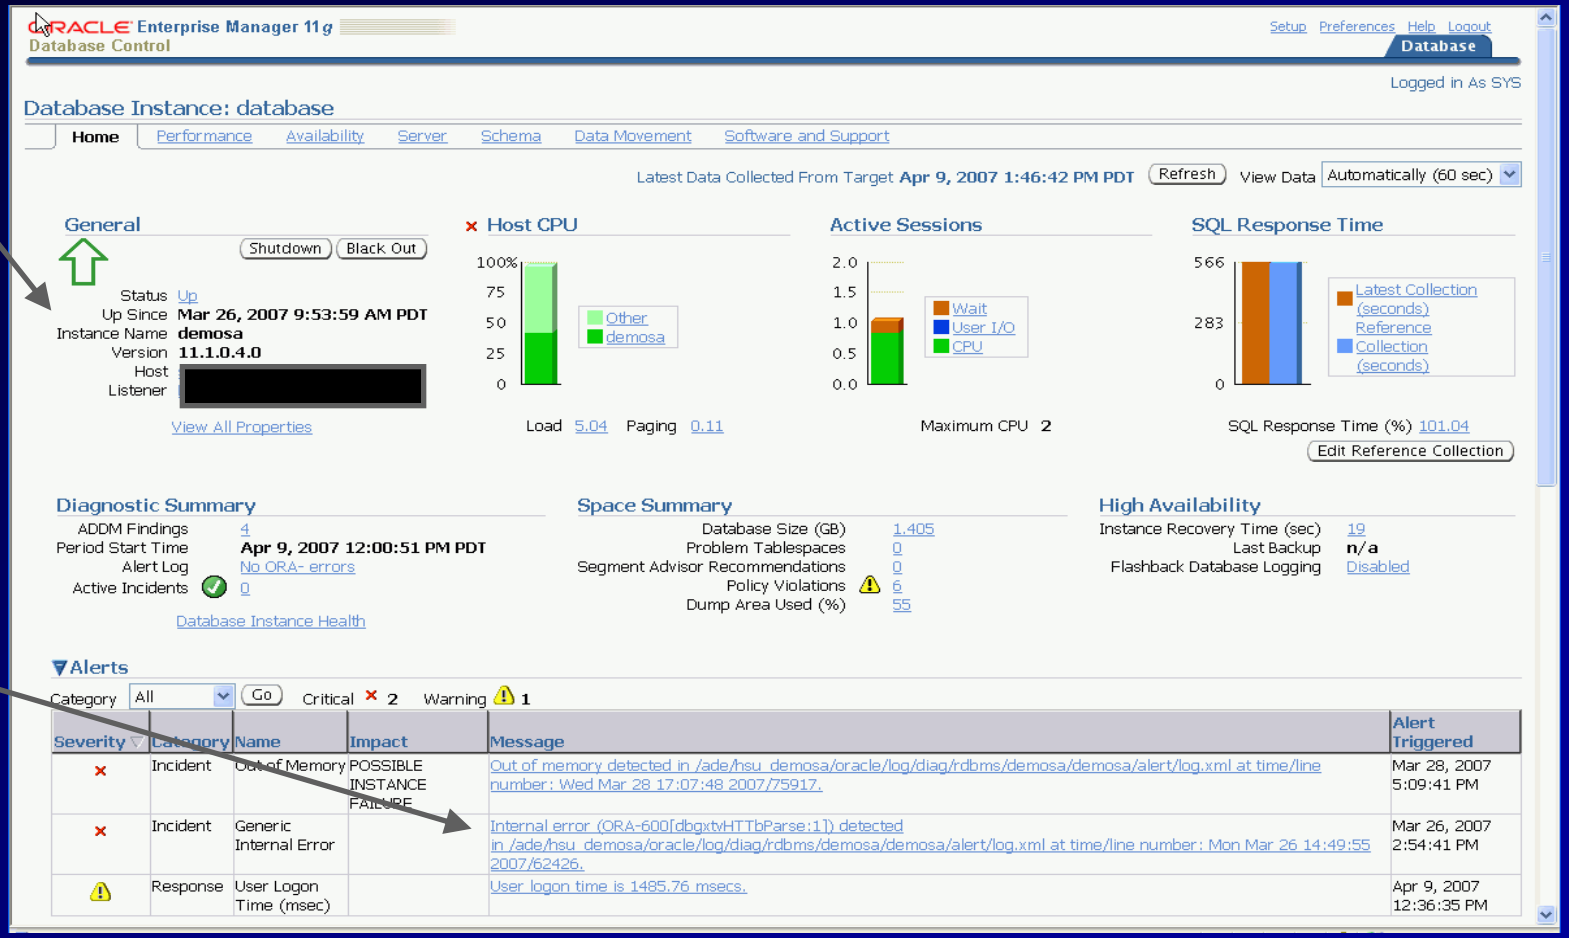

Support Workbench - Details

Click on SQL Repair Advisor

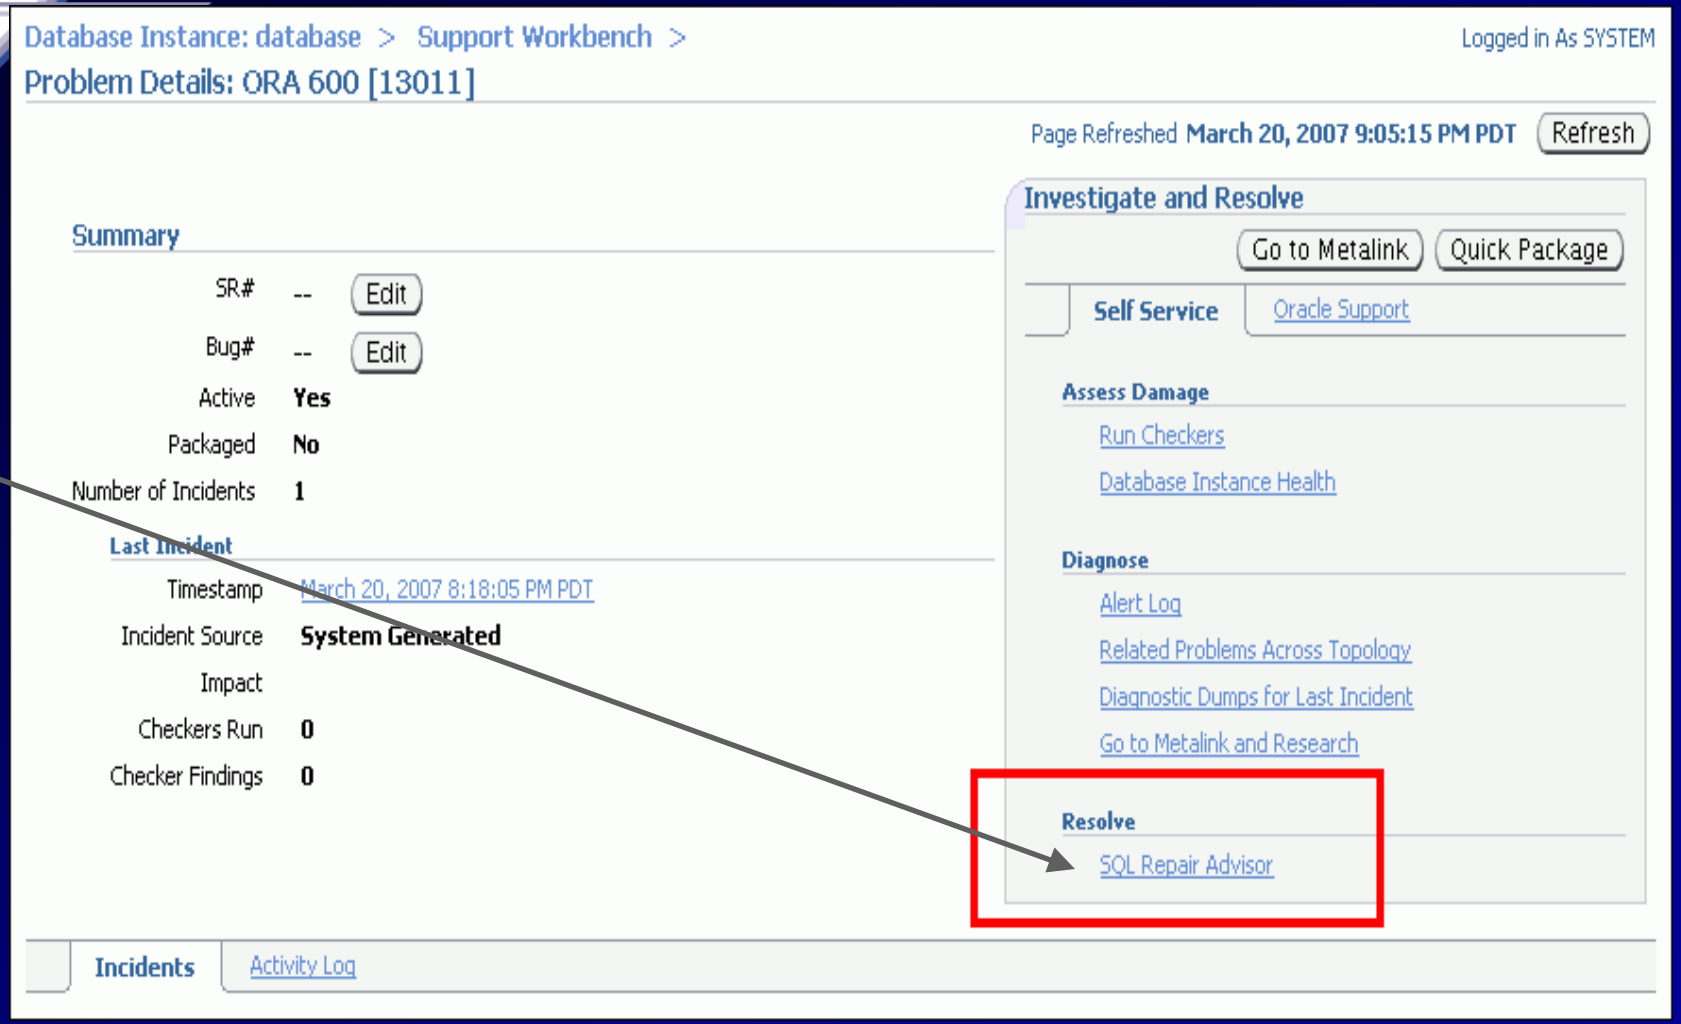

Results from SQL Repair Advisor

Click on View to Get the Detail finding of the Advisor

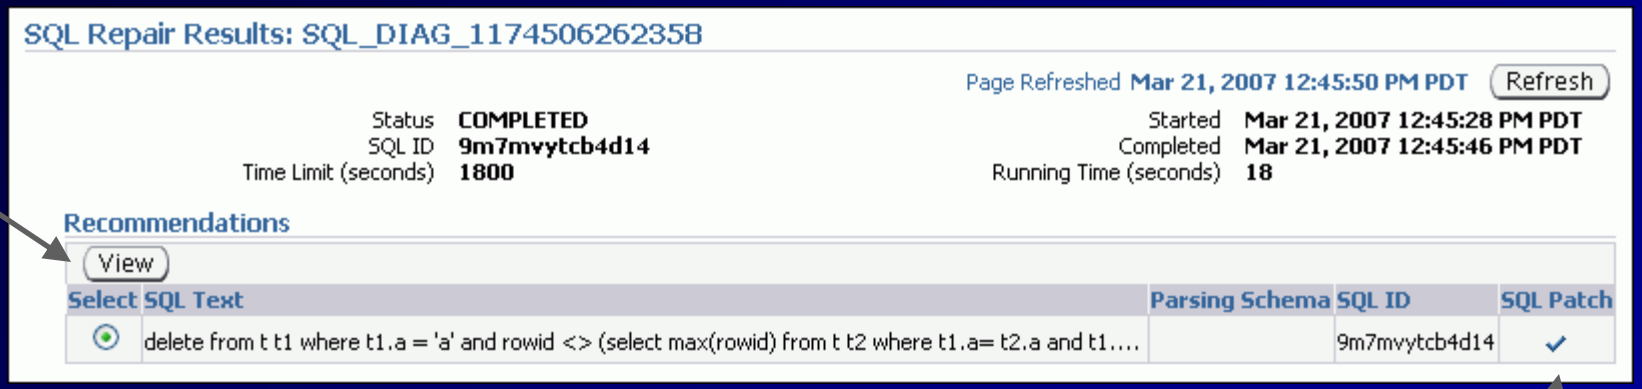

Note a SQL Patch (FIX for the SQL) has been generated

SQL Repair Advisor
Recommendation / Confirmation

Click on Implement To accept the SQL Patch

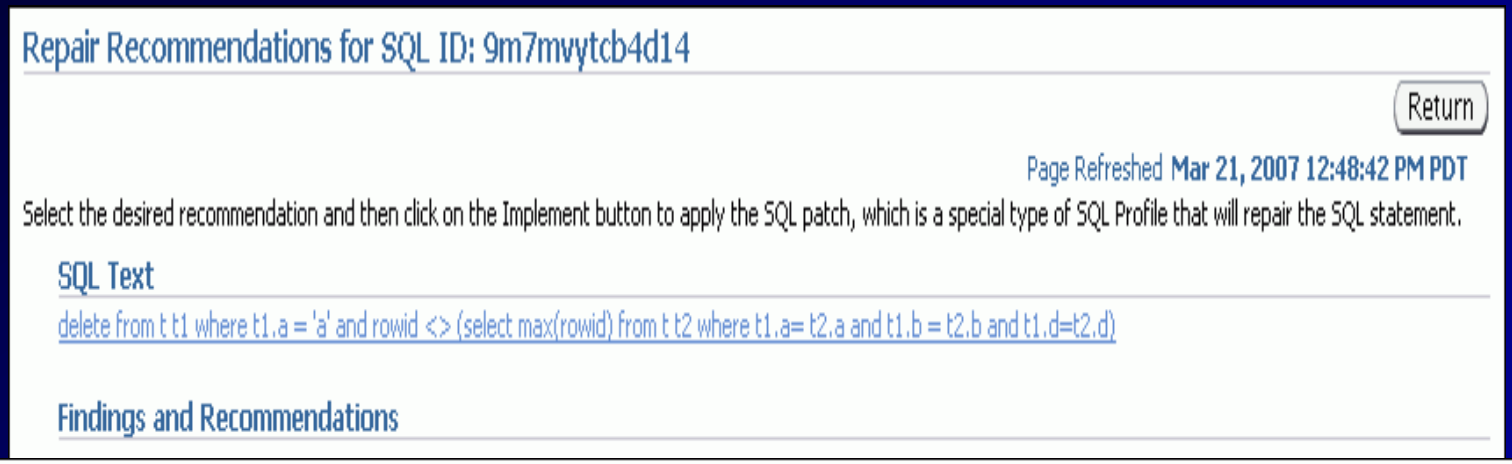

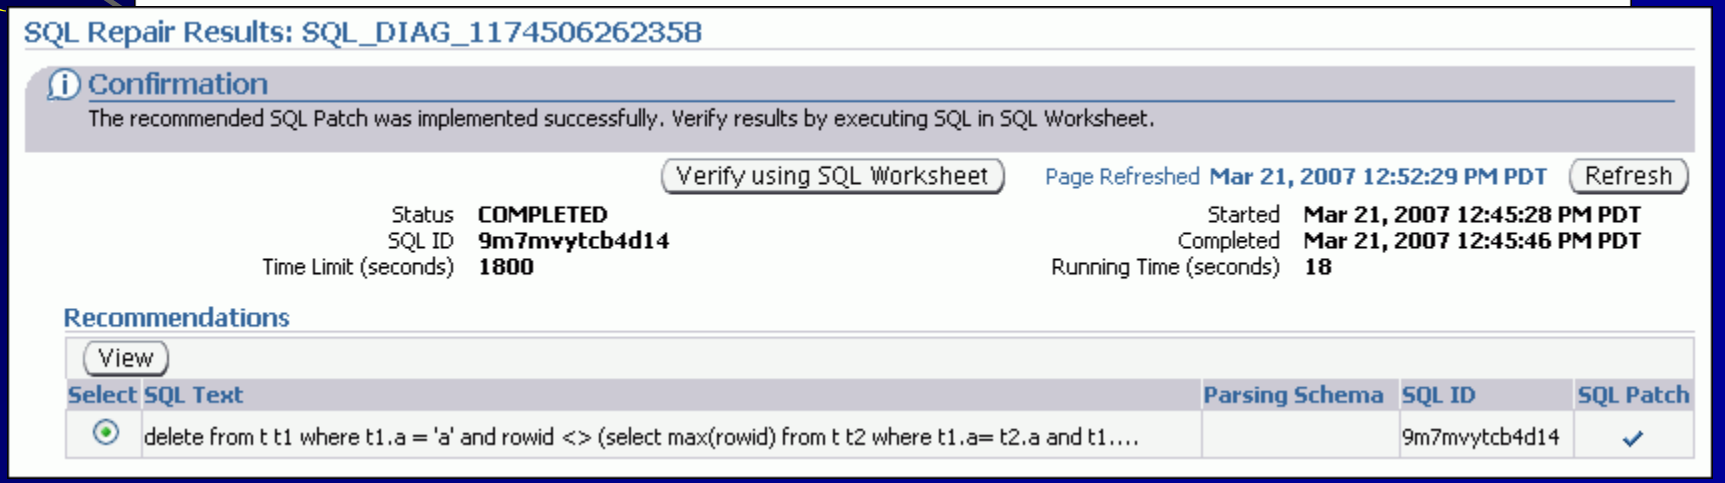

Adding Nodes the Easy Way in 11gR2 (FYI look only)

DB Instance – Software/Support
11gR2 - Deployment Procedures

Easy RAC Add Nodes (provisioning)

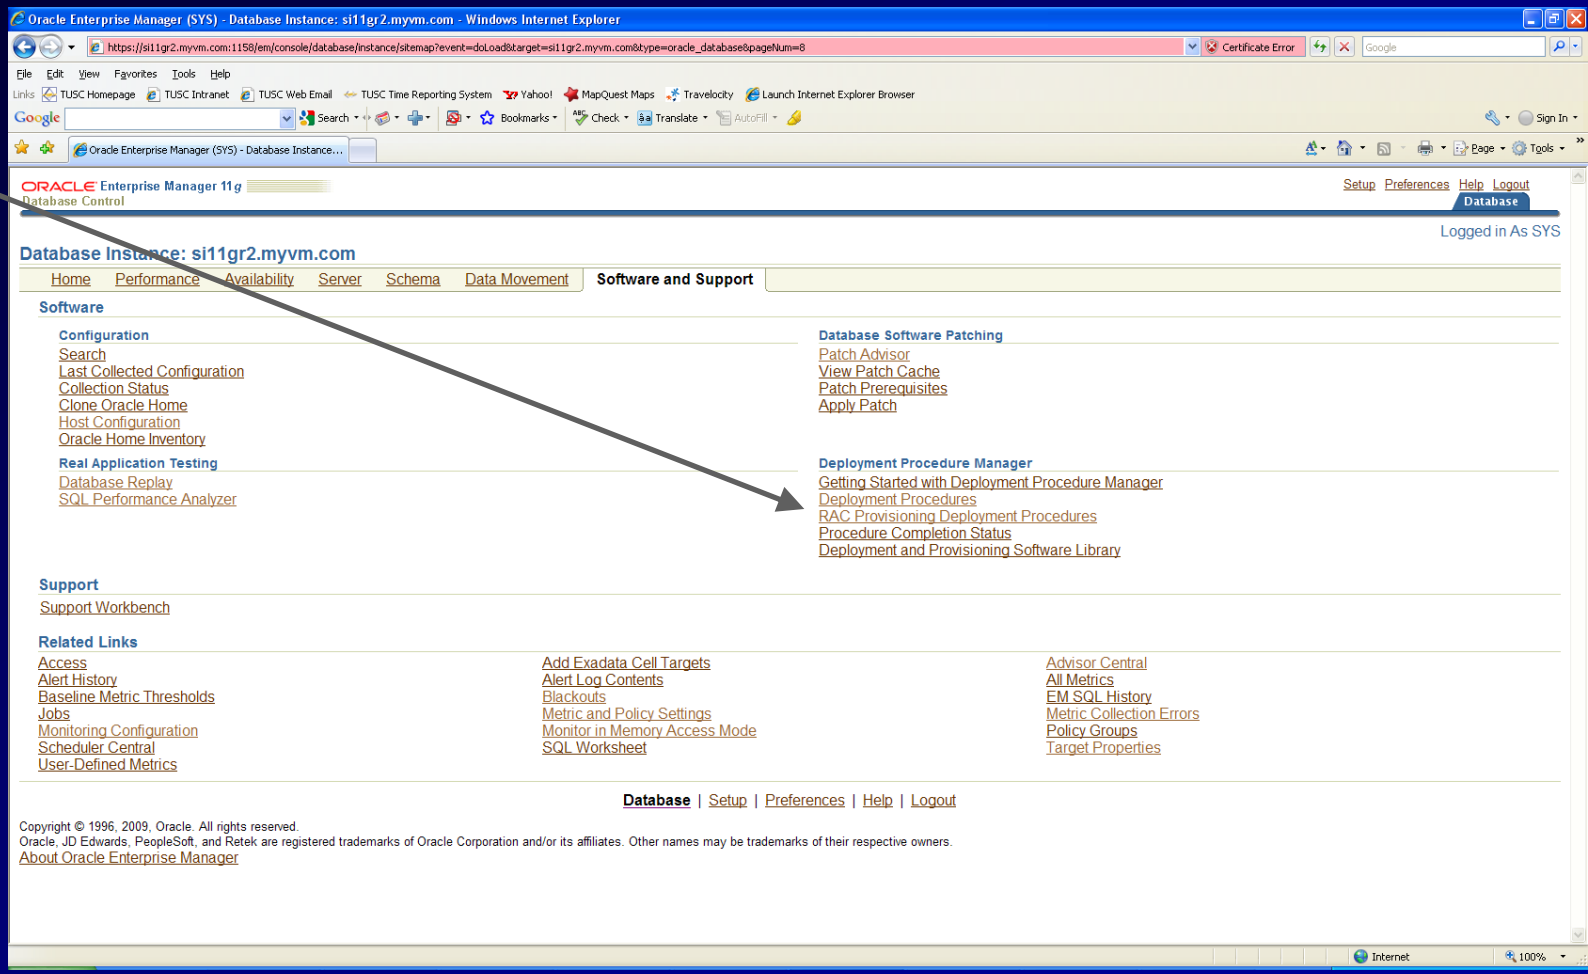

DB Instance – Software/Support
11gR2 - Deployment Procedures

One Click Extend Cluster

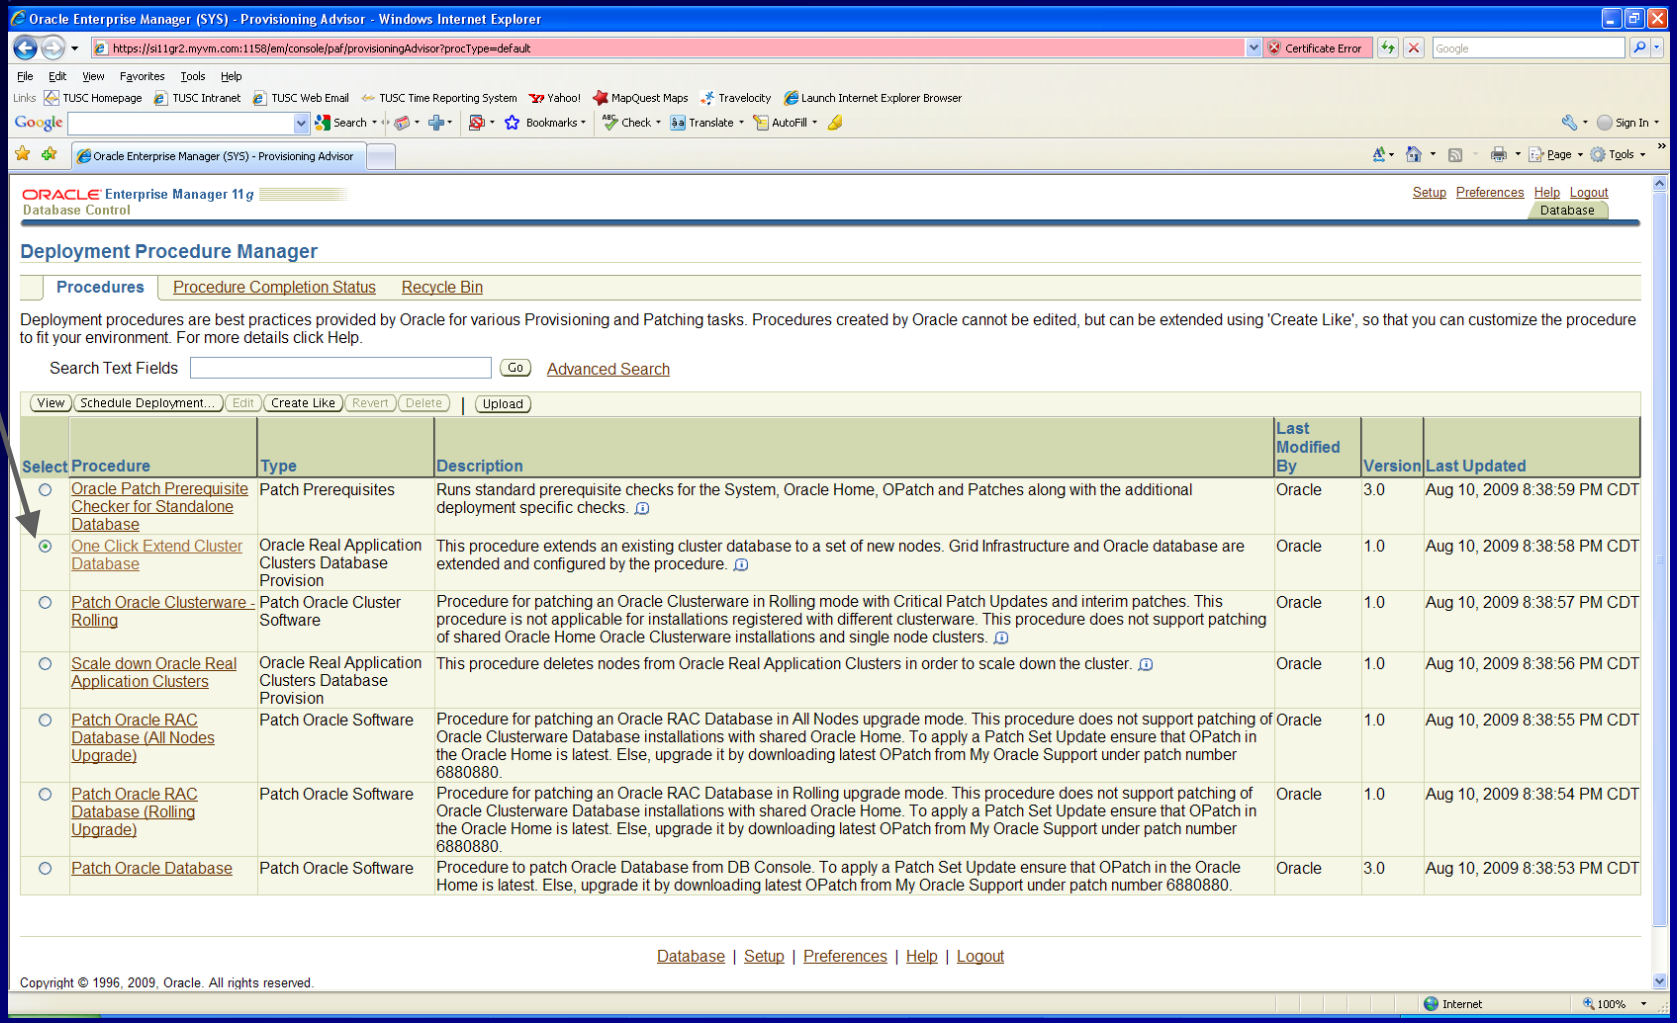

Helpful V$/X$ Queries (FYI Only)

V$ Views over the years

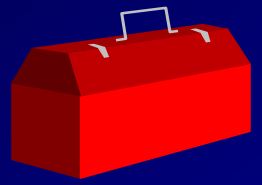

Version	V$ Views	X$ Tables
6	23	? (35)
7	72	126
8.0	132	200
8.1	185	271
9.0	227	352
9.2	259	394
10.1.0.2	340 (+31%)	543 (+38%)
10.2.0.1	396	613
11.1.0.6.0	**484 (+23%)**	**798 (+30%)**
11.2.0.1.0	**496 (+25%)**	**945 (+54%)**

Listing of V$ Views

```
select      name
from        v$fixed_table
where       name like 'GV%'
order by    name;

NAME
--------------------------------
GV$ACCESS
GV$ACTIVE_INSTANCES
GV$ACTIVE_SESS_POOL_MTH
GV$AQ1
GV$ARCHIVE...
```

Need GV$ - Instance ID

```
select (1 - (sum(decode(name, 'physical reads',value,0)) /
        (sum(decode(name, 'db block gets',value,0)) +
        sum(decode(name, 'consistent gets',value,0))))) * 100
    "Hit Ratio"
from     v$sysstat;

Hit Ratio
----------
90.5817699
```

Need GV$ - Instance ID

```
select    inst_id, (1 - (sum(decode(name, 'physical
    reads',value,0)) /
          (sum(decode(name, 'db block gets',value,0)) +
          sum(decode(name, 'consistent gets',value,0))))) * 100
    "Hit Ratio"
from      gv$sysstat
group by inst_id;

INST_ID      Hit Ratio
-----------  ----------------
         1  90.5817699
```

X$ used to create V$

```
select   *
from     v$fixed_view_definition
where    view_name = 'GV$INDEXED_FIXED_COLUMN';
```

VIEW_NAME	VIEW_DEFINITION
gv$indexed_fixed_column kqftanam, kqfconam,kqfcoipo	select c.inst_id, kqfcoidx, from X$kqfco c, X$kqfta t where t.indx = c.kqfcotab and kqfcoidx != 0

Listing of X$ Tables

```
select      name
from        v$fixed_table
where       name like 'X%'
order by    name;

NAME
--------------------------------
X$ACTIVECKPT
X$BH
X$BUFQM
X$CKPTBUF
X$CLASS_STAT...
```

Listing of X$ Indexes

(498 in 11.1.0.3.0, 419 in 10g; 326 in 9


```
select      table_name, index_number,
   column_name
from        gv$indexed_fixed_column
order by    table_name, index_number,
   column_name,                    column_position;
```

```
TABLE_NAME    INDEX_NUMBER COLUMN_NAME
----------------------------- ------------------------------ ------------------------------
   X$CLASS_STAT                        1 ADDR
   X$CLASS_STAT                        2 INDX
   X$DUAL                              1 ADDR
```

V$ - System Information

*select * from v$version;*

BANNER

Oracle Database 11g Enterprise Edition Release 11.1.0.3.0 - Beta

PL/SQL Release 11.1.0.3.0 - Beta

CORE 11.1.0.3.0 Beta

TNS for Linux: Version 11.1.0.3.0 - Beta

NLSRTL Version 11.1.0.3.0 - Beta

V$ - System Information


```
select *
from   v$option;

PARAMETER                      VALUE
------------------------------ ----------
   Partitioning                    TRUE
   Objects                         TRUE
   Real Application Clusters       FALSE
   Advanced Replication            TRUE
   Bit-Mapped Indexes              TRUE
```

...

V$ - V$SESSION_WAIT (waiting right now)


```
select event, sum(decode(wait_time,0,1,0)) "Waiting Now",
  sum(decode(wait_time,0,0,1)) "Previous Waits",
  count(*) "Total"
from v$session_wait
group by event
order by count(*);
```

WAIT_TIME = 0 means that it's waiting

WAIT_TIME > 0 means that it previously waited this many ms

V$ - V$SESSION_WAIT

EVENT	Waiting Now	Previous Waits	Total
db file sequential read	0	1	1
db file scattered read	2	0	2
latch free	0	1	1
enqueue	2	0	2
SQL*Net message from client	0	254	480

...

V$SESSION_WAIT
Current Specific waits

Buffer Busy Waits or Write Complete Waits Events:

```
SELECT     /*+ ordered */ sid, event, owner,
   segment_name,
              segment_type,p1,p2,p3
FROM       v$session_wait sw, dba_extents de
WHERE      de.file_id = sw.p1
AND                sw.p2 between de.block_id
              and de.block_id+de.blocks - 1
AND                (event = 'buffer busy waits'
                   OR event = 'write complete waits')
AND                p1 IS NOT null
ORDER BY event,sid;
```

V$EVENT_NAME Finding P1, P2, P3

```
col name for a20
col p1 for a10
col p2 for a10
col p3 for a10
select event#,name,parameter1 p1,parameter2 p2,parameter3
   p3
from  v$event_name
where name in ('buffer busy waits', 'write complete waits')
 EVENT# NAME                      P1         P2         P3
------------ ------------------          ---------- ---------- ----------
          74 write complete waits  file#      block#
```

V$ - V$SESSION_WAIT_HISTORY (Last 10 waits for session)

Buffer Busy Waits or Write Complete Waits Events:

```
SELECT    /*+ ordered */ sid, event, owner,
   segment_name,
          segment_type,p1,p2,p3
FROM      v$session_wait_history sw, dba_extents de
WHERE     de.file_id = sw.p1
AND               sw.p2 between de.block_id
          and de.block_id+de.blocks - 1
AND               (event = 'buffer busy waits'
                  OR event = 'write complete waits')
AND               p1 IS NOT null
ORDER BY event,sid;
```

Great V$ - V$SESSION_EVENT (waiting since the session started)

```
select sid, event, total_waits, time_waited,
    event_id
from v$session_event
where time_waited > 0
order by time_waited;
```

```
SID        EVENT                          TOTAL_WAITS TIME_WAITED
---------- ------------------------------ ----------- -----------
       159 process startup                          2           1
       167 latch: redo allocation                   4           1
       168 log buffer space                         2           3
       166 control file single write                5           4
...
```

V$ - V$SESSION_WAIT_CLASS (session waits by WAIT CLASS)

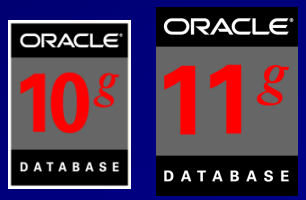

select wait_class, total_waits
from v$system_wait_class;

```
WAIT_CLASS           TOTAL_WAITS
-------------------- -----------
Other                       4180
Application                45269
Configuration                297
Concurrency                25467
Commit                     54805
Idle                     6925277
Network                  1859009
User I/O                  809979
System I/O               1103539
Scheduler                  10276
```

Great V$ - V$SYSTEM_EVENT (waits since the instance started)


```
select  event, total_waits, time_waited
from  v$system_event
where time_waited > 0
order by time_waited;
```

```
EVENT                          TOTAL_WAITS TIME_WAITED
------------------------------ ----------- -----------
latch: session allocation             5644           1
latch: redo allocation                   4           1
latch: cache buffers chains              4           3
enq: TX - index contention               1           3
direct path write temp                  57           6
row cache lock                           1           7  ...
```

V$ - Top 10 as % of All

```
select sum(pct_bufgets) percent
from (select rank() over ( order by buffer_gets desc )          as
    rank_bufgets,
   to_char(100 * ratio_to_report(buffer_gets) over      (), '999.99')
    pct_bufgets
from   v$sqlarea )
where  rank_bufgets  <  11;

PERCENT
-----------
     97.07
```

V$ - What Users are doing...


```
select  a.sid, a.username, s.sql_text
from    v$session a, v$sqltext s
where   a.sql_address       = s.address
and     a.sql_hash_value = s.hash_value
order   by a.username, a.sid, s.piece;

SID  USERNAME        SQL_TEXT
---- ---------------   ---------------------------------
   11  PLSQL_USER   update s_employee set salary = 10000
    9  SYS                 select a.sid, a.username, s.sql_text
    9  SYS                 from v$session a, v$sqltext
    9  SYS                 where a.sql_address  = s.address
   (...partial output listing)
```

Great V$ - V$SEGMENT_STATISTICS

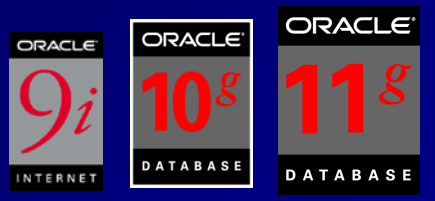

```
select object_name, statistic_name, value
from v$segment_statistics
where value > 100000
order by value;
```

OBJECT_NAME	STATISTIC_NAME	VALUE
ORDERS	space allocated	96551
ORDERS	space allocated	134181
ORDERS	logical reads	140976
ORDER_LINES	db block changes	183600

AWR – Segments by Buffer Busy Waits

Segments by Buffer Busy Waits

- % of Capture shows % of Buffer Busy Waits for each top segment compared
- with total Buffer Busy Waits for all segments captured by the Snapshot

Owner	Tablespace Name	Object Name	Subobject Name	Obj. Type	Buffer Busy Waits	% of Capture
				TABLE PARTITION	1,243,890	89.24
				INDEX PARTITION	39,258	2.82
				INDEX PARTITION	33,780	2.42
				INDEX PARTITION	28,563	2.05
				INDEX PARTITION	26,637	1.91

Back to Segment Statistics
Back to Top

Segments by Global Cache Buffer Busy

- % of Capture shows % of GC Buffer Busy for each top segment compared
- with GC Buffer Busy for all segments captured by the Snapshot

Owner	Tablespace Name	Object Name	Subobject Name	Obj. Type	GC Buffer Busy	% of Capture
SYS	SYSTEM	UNDO$		TABLE	21	100.00

Back to Segment Statistics

AWR – Segments by Logical Reads

Segments by Logical Reads

- Total Logical Reads: 225,112,503
- Captured Segments account for 90.9% of Total

Owner	Tablespace Name	Object Name	Subobject Name	Obj. Type	Logical Reads	%Total
				INDEX PARTITION	59,714,336	26.53
				INDEX PARTITION	48,936,144	21.74
				INDEX PARTITION	40,319,312	17.91
				INDEX PARTITION	23,612,368	10.49
				INDEX PARTITION	20,381,072	9.05

Back to Segment Statistics
Back to Top

Segments by Physical Reads

- Total Physical Reads: 1,722
- Captured Segments account for 41.1% of Total

Owner	Tablespace Name	Object Name	Subobject Name	Obj. Type	Physical Reads	%Total
				INDEX PARTITION	275	15.97
				INDEX PARTITION	164	9.52
				INDEX PARTITION	97	5.63

If Time Permits… the Future!

64-Bit advancement of Directly addressable memory

	Address Direct	Indirect/Extended
• 4 Bit:	16	(640)
• 8 Bit:	256	(65,536)
• 16 Bit:	65,536	(1,048,576)
• 32 Bit:	4,294,967,296	
• 64 Bit:	18,446,744,073,709,551,616	

- When the hardware physically implements the theoretical possibilities of 64-Bit, things will dramatically change….
  …moving from 32 bit to 64 bit will be like moving from 4 bit to 32 bit or like moving from 1971 to 2000 overnight.

64bit allows Directly Addressing 16 Exabytes of Memory

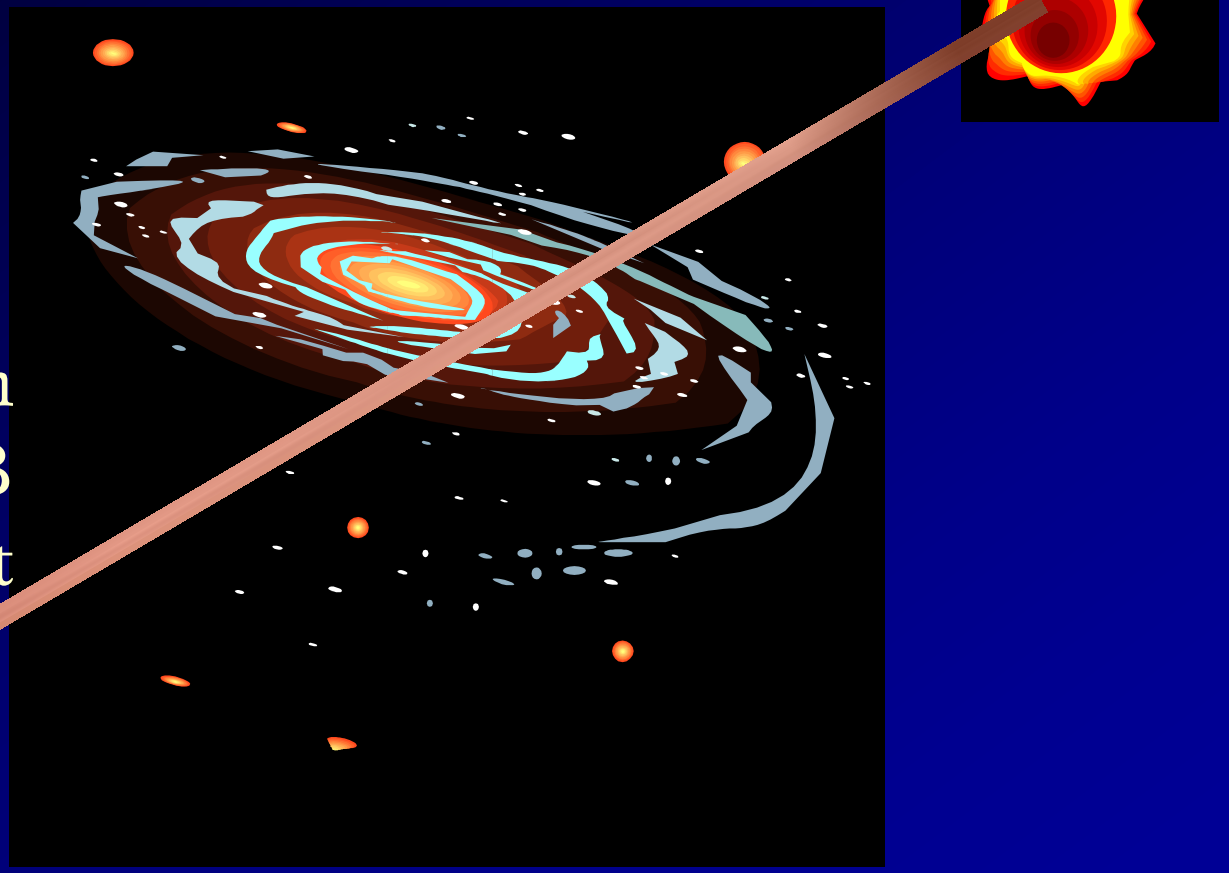

Stack single sheets (2K worth of text on each) about 4.8B miles high to get 16E!!

You could stack documents from the Earth so high they would pass Pluto!

The Future: 8 Exabytes
Look what fits in one 10g Database!

2K – A typewritten page
5M – The complete works of Shakespeare
10M – One minute of high fidelity sound
2T – Information generated on YouTube in one day
10T – 530,000,000 miles of bookshelves at the Library of Congress
20P – All hard-disk drives in 1995 (or your database in 2010)
700P –Data of 700,000 companies with Revenues less than $200M
1E – Combined Fortune 1000 company databases (average 1P each)
1E –Next 9000 world company databases (average 100T each)
8E – Capacity of ONE Oracle10g Database (CURRENT)
12E to 16E – Info generated before 1999 (memory resident in 64-bit)
16E – Addressable memory with 64-bit (CURRENT)
161E – New information in 2006 (mostly images not stored in DB)
1Z – 1000E (Zettabyte - Grains of sand on beaches -125 Oracle DBs)
100TY - 100T-Yottabytes – Addressable memory 128-bit (FUTURE)

What we covered:

- Kramer's missing Backup
- George's Untuned System
- Elaine's Untouched System
- Jerry's Perfect Tuning Plan
- Statspack / AWR
  - Top Waits
  - Load Profile
  - Latch Waits
  - Top SQL
  - Instance Activity
  - File I/O
- The Future EM & ADDM
- Helpful V$/X$
- Summary

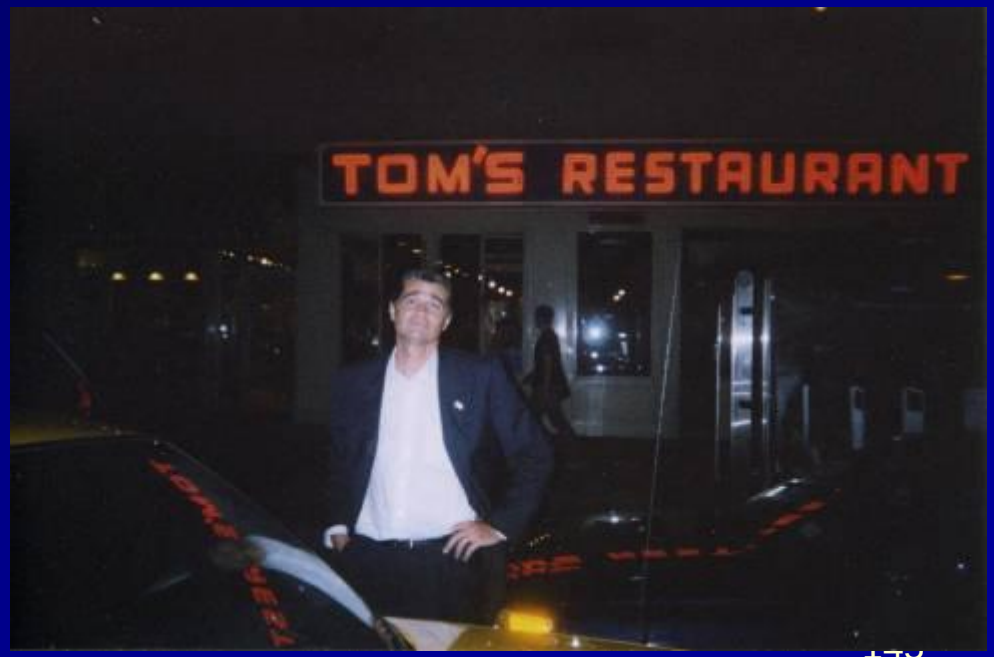

Questions??

"Perfection is achieved, not when there is nothing left to add, but when there is nothing left to take away."

--Antoine de St. Exupery

For More Information

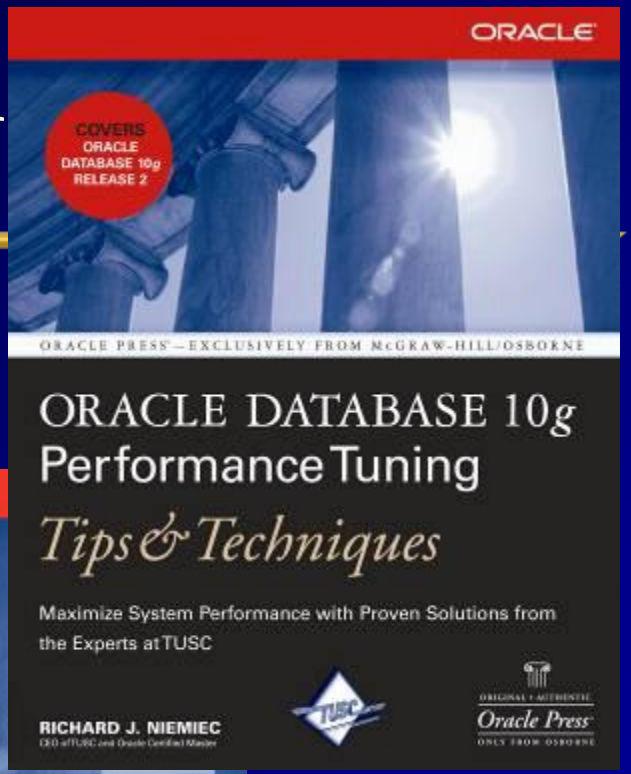

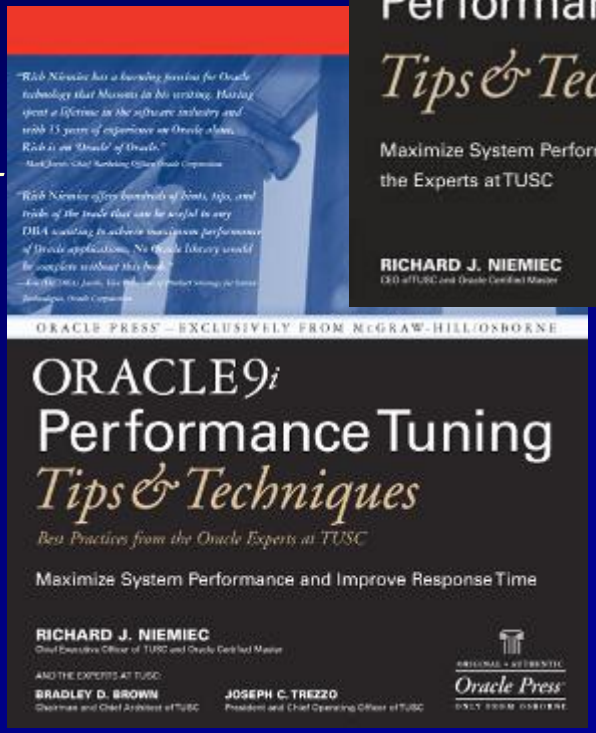

- www.tusc.com
- *Oracle9i Performance Tuning Tips & Techniques; Richard J. Niemiec; Oracle Press (May 2003)*
- *Oracle 10g Tuning (June 11, 2007)*

"If you are going through hell, keep going" - Churchill

References

- www.tusc.com, www.oracle.com, www.ixora.com, www.laoug.org, www.ioug.org, technet.oracle.com
- Oracle 9i, 10g, 11g documentation
- *Oracle10g Performance Tuning Tips & Techniques*, Rich Niemiec
- *Oracle PL/SQL Tips and Techniques, Joseph P. Trezzo; Oracle Press*
- *Oracle9i Web Development, Bradley D. Brown; Oracle Press*
- *Special thanks to Steve Adams, Mike Ault, Brad Brown, Kevin Gilpin, Herve Lejeune, Randy Swanson and Joe*

References

- *The Self-managing Database: Automatic Performance Diagnosis;* Karl Dias & Mark Ramacher, Oracle Corporation
- *EM Grid Control 10g;* otn.oracle.com, Oracle Corporation
- Oracle Database 10g *A*utomated Features , Mike Ault, TUSC
- *Oracle Enterprise Manager 10g: Making the Grid a Reality*; Jay Rossiter, Oracle Corporation
- *The Self-Managing Database: Guided Application and SQL Tuning;* Benoit Dageville, Oracle Corporation
- *The New Enterprise Manager: End to End Performance Management of Oracle;* Julie Wong & Arsalan Farooq

References

- www.tusc.com. www.rolta.com
- *Oracle10g Performance Tuning Tips & Techniques; Richard J. Niemiec; Oracle Press*
- Database Secure Configuration Initiative: Enhancements with Oracle Database 11g, www.oracle.com
- All Oracle11g Documentation from Oracle Beta Site
- Introduction to Oracle Database 11g, Ken Jacobs
- Oracle Database 11g New Features, Linda Smith
- New Optimizer Features in 11g, Maria Colgan
- www.ioug.org, www.oracle.com, en.wikipedia.org & technet.oracle.com
- Thanks Dan M., Bob T., Brad, Joe, Heidi, Mike K., Debbie, Maria, Linda
- All companies and product names are trademarks or registered trademarks of the respective owners.

Rolta TUSC – *Your* Partner …. Accomplished in Oracle!

2010 Oracle Partner of the Year (7 Titans Total)

Prior Years Winner 2002, 2004*, 2007*, 2008

*Won 2 Awards

Rolta TUSC Services

- **Oracle**
  - E-Business Suite implementation, R12 upgrades, migration & support
  - Fusion Middleware and Open Systems development
  - Business Intelligence (OBIEE) development
  - Hyperion Financial Performance Management
  - DBA and Database tactical services
  - Strategic Global Sourcing
- **IT Infrastructure**
  - IT Roadmap - Security & Compliance - Infrastructure Management
  - Enterprise Integration / SOA - High Availability and Disaster Planning
- **Profitability & Cost Management**
  - Financial Consolidation - Budgeting & Forecasting
  - Profitability & Risk Analysis - Enterprise Performance Management
  - Operational, Financial & Management Reporting
- **Rolta Software Solutions**
  - *i*Perspective™ - rapid data & systems integration
  - Geospatial Fusion™ - spatial integration & visualization
  - OneView™ - business & operational intelligence

Copyright Information

- Neither Rolta TUSC nor the author guarantee this document to be error-free. Please provide comments/questions to rich@tusc.com. I am always looking to improve!
- Rich Niemiec/ Rolta TUSC © 2011. This document cannot be reproduced without expressed written consent from Rich Niemiec or an officer of Rolta TUSC, but may be reproduced or copied for presentation/conference use.

Contact Information

Rich Niemiec: rich@tusc.com

www.tusc.com

Rich's Overview (rich@tusc.com)

- Advisor to Rolta International Board
- Former President of TUSC
  - Inc. 500 Company (Fastest Growing 500 Private Companies)
  - 10 Offices in the United States (U.S.); Based in Chicago
  - Oracle Advantage Partner in Tech & Applications
- Former President Rolta TUSC & President Rolta EICT International
- Author (3 Oracle Best Sellers – #1 Oracle Tuning Book for a Decade):
  - Oracle Performing Tips & Techniques (Covers Oracle7 & 8i)
  - Oracle9i Performance Tips & Techniques
  - Oracle Database 10g Performance Tips & Techniques
- Former President of the International Oracle Users Group
- Current President of the Midwest Oracle Users Group
- Chicago Entrepreneur Hall of Fame - 1998
- E&Y Entrepreneur of the Year & National Hall of Fame - 2001
- IOUG Top Speaker in 1991, 1994, 1997, 2001, 2006, 2007
- MOUG Top Speaker Twelve Times
- National Trio Achiever award - 2006
- Oracle Certified Master & Oracle Ace Director