

◦ Sherlock Holmes for DBA's



Kellyn Pot'Vin
Sr. Technical Consultant
Enkitech

Denver DBA

- Kellyn Pot'Vin,
- Westminster, Colorado
- Oracle ACE, Sr. Technical Specialist at Enkitech
- Multi-Platform DBA
- Specialize in performance and management of large enterprise environments.
- Board of directors for RMOUG, Director of Training Days 2013, KSCOPE DB Track Lead 2013
- Blog: DBAKevlar.com
- Twitter: [@DBAKevlar](https://twitter.com/DBAKevlar)

Challenges for Today's DBA

- Added Complexity by Less Manual Configurations in Environment.
- Added Roles to DBA's, either requiring more specializing and/or overall knowledge required.
- Added Features, systems.
- More open source solutions required, outside of standard design.

The Pitfall of All Optimization: The “Tinsel Monkey”



“Tinsel” vs. “Valuable”

- Tinsel = anything that is shiny-
 - graphs from EM without investigation.
 - costs without research verifying are valid.
 - Assumptions without proof.
- Tuning “Tinsel” results in no productive gain in performance in the database.
- Valuable optimization=Performance Gain.
- Performance gain should equal time savings.
- Time is gold.

What is Performance “Gold”

- Performance optimization that-
 - saves valuable time to a long running process.
 - Saves users valuable time to their work day.
 - Saves results in value to the business.
 - Re-allocates valuable resources, (memory, I/O or sessions) back to the database.
- Optimizing that saves time from maintenance work that can be re-allocated back to the business processing.

Saving DBA Time

- Sharing knowledge- Optimization should never be viewed as a “mystery process” or “voodoo”.
- Developers and Analysts should have access to DBA Performance Views.
- Ability to create explain plans.
- Access to Enterprise Manager “Top Activity” and drill downs.

Goals for the Attendee

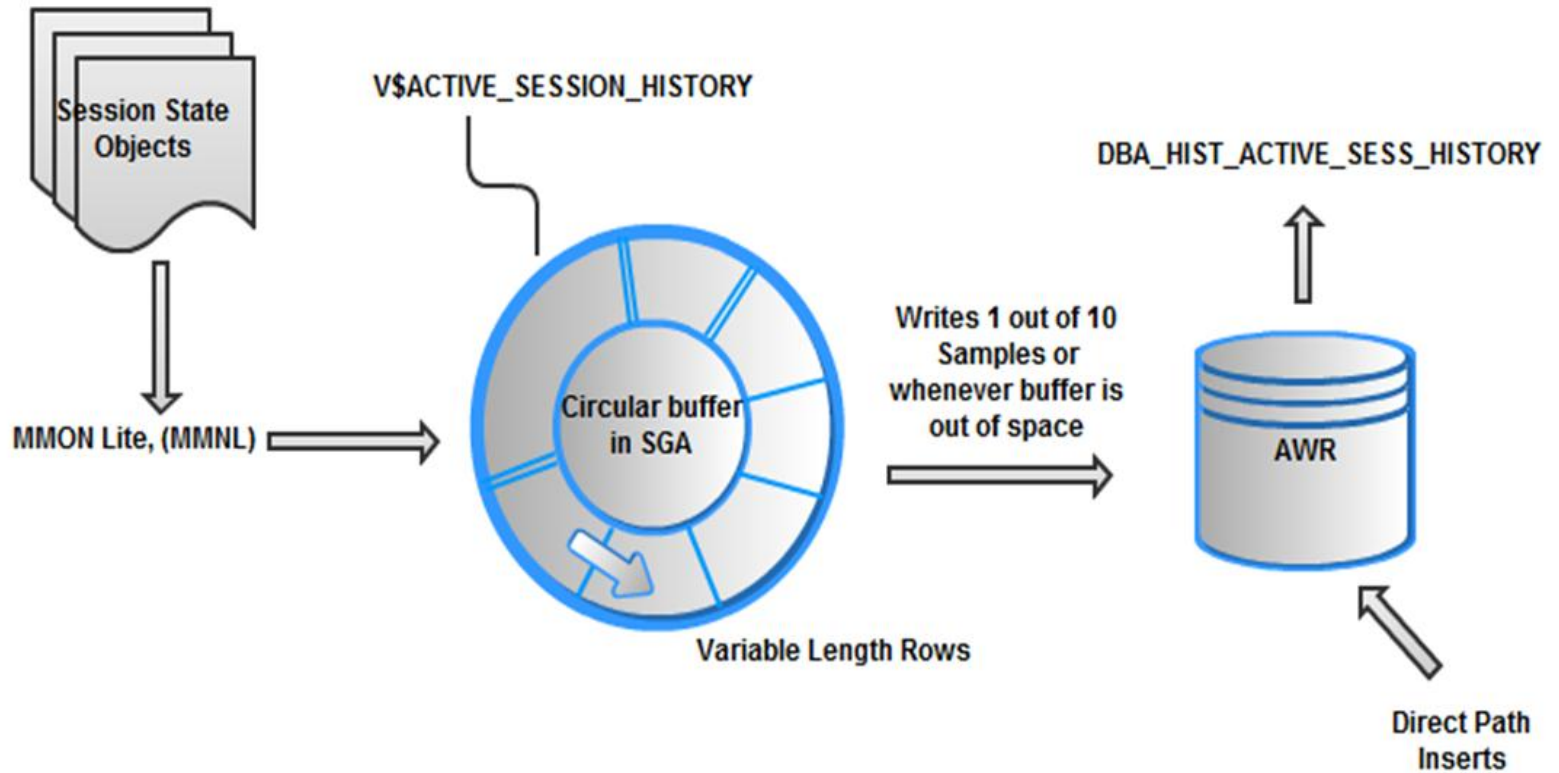
- AWR, ASH and Tracing Knowledge.
- Steps to trouble-shooting an issue.

Over all, the goal of this session is for you to always, **ALWAYS, ALWAYS** base any changes you make on data and proof. **NEVER** base changes on assumptions, guesses or because you “read it somewhere”.

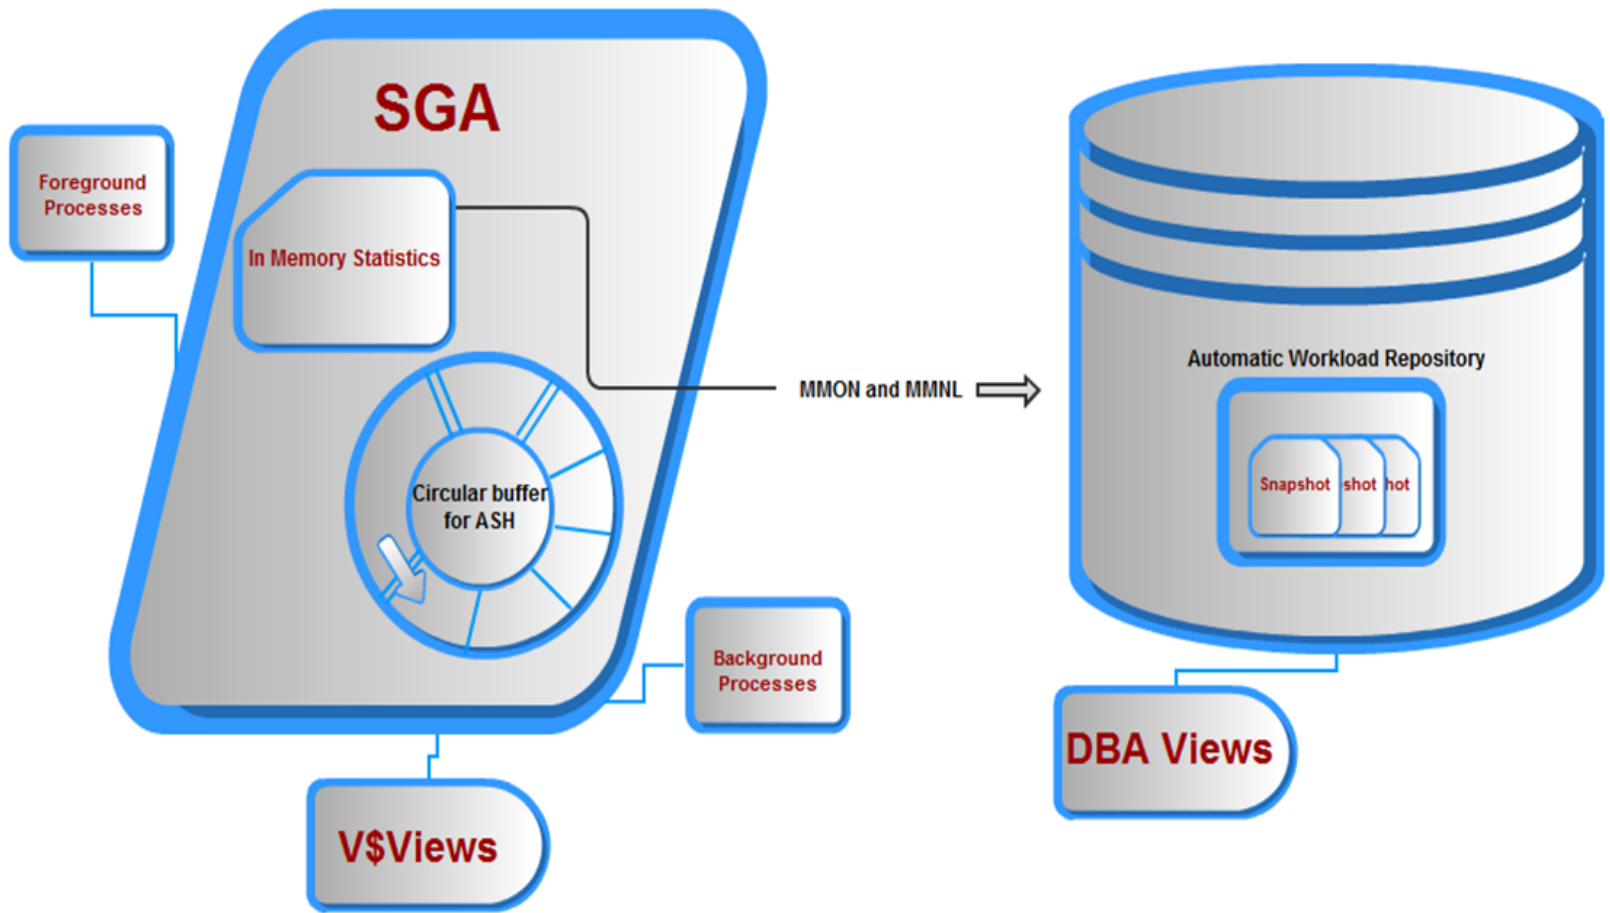
Presentation Agenda

- Architecture
- History
- Reports
 - ADDM
 - ASH
 - AWR
 - Explain Plans
- Trace Files
 - I0046
 - I0053

ASH Architecture



AWR Architecture



AWR Repository

- Used not only by the AWR reports
 - Automatic Database Diagnostic Monitor, (ADDM Reporting)
 - SQL Tuning Advisor
 - Segment Advisor
- By default, snapshots every hour. Retention is for 7 days. Both are modifiable.
- Snapshots can be taken at any time:

```
EXEC DBMS_WORKLOAD_REPOSITORY.create_snapshot;
```

Running Tuning Pack Reports

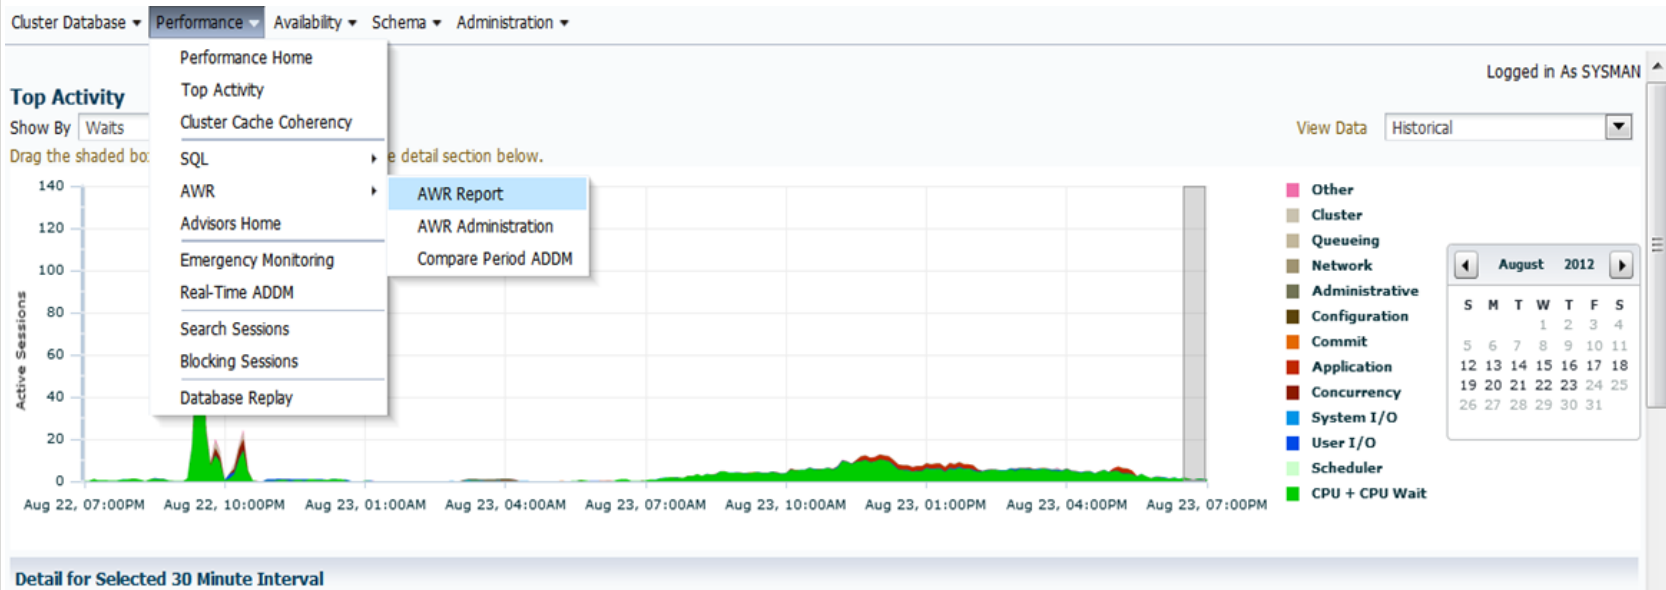
- Through Enterprise Manager
- `$ORACLE_HOME/rdbms/admin`
 - ADDM- `addmrpt.sql`
 - ASH- `ashrpt.sql`
 - AWR- `awrrpt.sql`
 - SQL_ID Specific AWR- `awrsqrpt.sql`
- These reports can be in HTML or TEXT versions.
- Require the Tuning and Management Pack license.

AWR Reports

Statspack on “Steroids”

- As detailed snapshot of database depending on intervals of snapshot setup, (hourly be default.)
- Introduced in Oracle 10g
- Evolution to statspack, requests for performance reporting improvements.
- “Always on” approach to performance metrics with requirement of non-locking collection process.
- Requires Management Tuning Pack License from Oracle.

Running AWR from Enterprise Manager



- Post logging into the database, click on Performance.
- From Performance → AWR → AWR Report

ASH Report

- Best in smaller timeslots.
- Includes `SQL_ID`, object information and timeline of waits.
- Samples each active database session every second.
- Data is held in buffer in memory.
- Built into the Oracle kernel and accessed through the `v$active_session_history` view.
- In an AWR snapshot, 1 row in IO from ASH buffer is placed into the AWR repository.
- Managed by the MMNL, (Memory Monitor Lite)
- Should not be used to track occurrence.

Running ASH Report from Enterprise Manager

- ASH is always by time, not snapshot.
- Set start date and time.
- End date and time
- Generate report

Oracle Database ▾ Performance ▾ Availability ▾ Schema ▾ Administration ▾

Logged in As SYSMAN

Run ASH Report

Specify the time period for the report.

Start Date 
(Example: 12/15/03)

End Date 
(Example: 12/15/03)

Start Time AM PM

End Time AM PM

Filter

HTML Format ASH Report

ASH Report For

(1 Report Target Specified)

DB Name	DB Id	Instance	Inst num	Release	RAC	Host
	2601412324		2	10.2.0.5.0	YES	om

CPUs	SGA Size	Buffer Cache	Shared Pool	ASH Buffer Size
24	65,536M (100%)	55,280M (84.4%)	10,279M (15.7%)	40.5M (0.1%)

	Sample Time	Data Source
Analysis Begin Time:	28-Aug-12 14:49:15	V\$ACTIVE_SESSION_HISTORY
Analysis End Time:	28-Aug-12 14:54:15	V\$ACTIVE_SESSION_HISTORY
Elapsed Time:	5.0 (mins)	
Sample Count:	299	
Average Active Sessions:	1.00	
Avg. Active Session per CPU:	0.04	
Report Target:	SQL_ID like 'bv1nfusy2nxkd'	18.9% of total database activity

ASH Report

- [Top Events](#)
- [Load Profile](#)
- [Top SQL](#)
- [Top PL/SQL](#)
- [Top Sessions](#)
- [Top Objects/Files/Latches](#)
- [Activity Over Time](#)

[Back to Top](#)

Top Events

- [Top User Events](#)
- [Top Background Events](#)
- [Top Event P1/P2/P3 Values](#)

[Back to Top](#)

ADDM Report High Lights

“Just the Facts” Report

- Produced from AWR Snapshots
- Broken down by “Findings and Recommendations”
- Identifies high-level misconfigurations.
- Lists what “Is Not an Impact”
- Do NOT recommend for non-DBA Personnel.

Creating Explain Plans

Creating an Explain plan:

```
EXPLAIN PLAN FOR <SQL STATEMENT>;  
SELECT PLAN_TABLE_OUTPUT from  
  table(dbms_xplan.display('plan_table',null,'basic'));
```

Preferred:

```
SELECT * FROM  
  TABLE(DBMS_XPLAN.DISPLAY);
```

From a previous run, using SQL_ID:

```
SELECT * FROM  
TABLE(DBMS_XPLAN.DISPLAY_AWR('SQL_ID'  
  D'));
```

Trace Events and Levels

Event	Level	Description
10046	0	No statistics generated
10046	1	Standard trace output including parsing, executes and fetches plus more
10046	2	Same as Level 1
10046	4	Level 1 + Bind Variables
10046	8	Level 1 + Waits
10046	12	Level 1 + Bind Variables & Waits
10046	16	Added in 11g to generate STAT line dumps for each execution. STAT dumping has been amended in 11g so that they are not aggregated across all executions but are dumped after execution.
10053	1	Dump Optimizer Statistics and Computations
10053	2	Dump Optimizer Computations only

Tracing: 10046

- **10046**

```
alter session set events '10046 trace name
    context forever, level 8';

exec
dbms_system.set_sql_trace_in_session(<sid>,<s
erial#>,TRUE);

exec
dbms_system.set_ev(<sid>,<serial#>,10046,12,'
');
```

Tracing: I 0053

```
Alter session set events '10053 trace  
name context forever, level 1';
```

OR

```
begin
```

```
DBMS_SQLDIAG.DUMP_TRACE (p_sql_id=> '<sql  
_id>',  
p_child_number=>0,  
p_component=>'Optimizer',  
p_file_id=> '<distinguish trace  
name>');
```

```
end;
```

```
/
```

Trace Files are Too Big/Complex

- Become familiar with TKPROF for 10046 trace files.
- If tracing often, sell cost effectiveness like advanced trace assistance tools like Method R's "Mr.Trace".
- Become confident in how to find bind variable, sql_id and column data values in a trace file.
- Become friends with find/search.. 😊

Real Life Scenario I

Developer notices that performance is poor in two standard transformation process.

Process utilizes CTAS, (create table as)

Believes parallel process is being degraded to serial due to limitations in CPU resources in database.

First Step?

Second Step?

Scenario 1, Step 1

- Explain plan of first CTAS shows that although SQL has parallel hints in “select” section, no parallel was designated in “create table” section of statement.
- Second process showed parallel in explain plan, but processing via Enterprise Manager did show only one process active for majority of time, other parallel slaves waiting.

Scenario 1, Step 2

- I0046, level 8 trace of second process.
- Trace file clearly shows single slave session busy on function call, where parallel at all other processing.
- Inspection of code showed function needed to be created with “parallel_enable” clause addition.

Real Life Scenario 2

11g Upgrade recently performed in environment.

Overall performance good, but number of large DML processes degrading.

Elapsed time extended, but time extension is spread across each transaction in the overall processing.

First step?

Second step?

Scenario 2, Step 1

- AWR reports show significant increase in times.
- IO has increased, physical reads on tables involved in DML shown as increase.
- No change in explain plans, utilizing 10g plans or new plans show no improvement.

Scenario 2, Step 2

- 10053 trace, level 1 by SQL_ID's of DML involved.
- Discovered in the trace file, full scans of tables due to following 11g parameter:
- `CELL_OFFLOAD_PROCESSING=TRUE`
- Per the MOS note- High 'ksv master wait' And 'ASM File Metadata Operation' Waits In Non-Exadata 11g [ID 1308282.1]

Real Life Scenario 3

- Users complain of performance degradation of main processing in data mart when larger, heavy concurrent processes.
- Same code executed in database when degradation occurs.
- AWR shows adequate SGA, PGA hit, high I/O.
- First step? Second Step?

Scenario 3, First Step:

- **Memory and PGA show fine, yet high I/O, inspect the memory ,workarea and temp:**

```
select vst.sql_text, swa.sql_id, swa.sid,  
       swa.tablespace  
       , swa.operation_type  
       , trunc(swa.work_area_size/1024/1024) "PGA MB"  
       , trunc(swa.max_mem_used/1024/1024) "Mem MB"  
       , trunc(swa.tempseg_size/1024/1024) "Temp MB"  
from v$sql_workarea_active swa, v$session vs,  
     v$sqltext vst  
where swa.sid=vs.sid  
and swa.sql_id=vs.sql_id  
and vs.sql_id=vst.sql_id  
and vst.piece=0  
order by swa.sql_id;
```


Scenario 3, Second Step:

- Workarea shows no PGA allocated, considering the size of temp table created as part of the CTAS, hash joins and hash group by sort.
- Performed 10048 trace of session.
- Action Plan: Traced session and noted the following, which was then located in procedure call: `alter session set WORKAREA_SIZE_POLICY = MANUAL;`

Optimization Data

- Database Level:
 - AWR Reports
 - ASH Reports
 - ADDM Reports
 - Enterprise Manager
 - Queries
- Session Level:
 - Explain Plans
 - Traces
 - Queries

Design of an Optimization Exercise

- Collect top 10 list of database performance challenges.
- Review with peers and user base.
- Create top 10 list that incorporates both the DBA list of top 10 and the User list.
- Divide the list into “low-hanging fruit” and those that will demand extensive resources to correct.
- Review the database again only after the “low-hanging fruit” has been addressed.

Summary

- Optimization is not an exact science, but following logical steps when scenarios develop is key.
- Always optimize for time savings
- Research to produce data, (proof) behind any optimize goaled exercise to eliminate wasted effort before an exercise is commenced.



Thank you

Kellyn Pot'Vin

Website: <http://dbakevlar.com>

and <http://enkitec.com>

Email: dbakevlar@gmail.com or kpotvin@enkitec.com

Twitter: @DBAKevlar

Facebook and LinkedIn: Kellyn Pot'Vin