Utah Oracle Users Group

SageLogix inc.
Oracle and e-Business Expertise

Optimizing The Performance of Oracle-based Applications

Tuning for DBAs

Tim Gorman
Tim@SageLogix.com

Principal - http://www.SageLogix.com

Agenda for Today

SageLogix inc.
Oracle and e-Business Expertise

- **Tuning for DBAs**
  - Overall Performance Tuning Methodology
  - Identifying problems from a global perspective
- **Tuning for Developers and DBAs**
  - Analyzing problems specifically
  - Tuning SQL statements

Agenda

SageLogix inc.
Oracle and e-Business Expertise

- Overall Performance Tuning Methodology
  - Gathering, understanding, and interpreting statistics
    - Empiricism
  - Prioritizing for maximum effect
    - Triage
- Identifying problems globally
  - Time measures in Oracle
    - Response-time
      - Service-Time
      - Wait-time

Common situations

SageLogix inc.
Oracle and e-Business Expertise

- Situation #1

  Application is an ERP product (Oracle Manufacturing). While applying a patch which updates data in the database, performance is so poor (40 hours) that 2 full days of downtime each are needed to apply the patch to each system (DEV, TEST, and PROD).

- Situation #2

  Overall performance is lousy for a partly-customized ERP system. Developers are trying to tune SQL as they detect problems, but it does not seem to be having much impact. Lack of performance is effectively depriving the business of expected functionality.

Case Studies

SageLogix inc.
Oracle and e-Business Expertise

- Both situations are *scary* due to lack of information
  - Phone- or web-based Oracle Support is ineffective
    - Support Analysts can only act on information given
    - Performance problems are often *unique* due to all the variables in the situation
  - Where to start looking for information?
    - Is it a hardware problem? Operating system? RDBMS? Network? Application? A bug?
  - Once there is some indication on the approximate causes, how do we fix it?
    - Quickly!
    - Effectively!
  - How do we keep it from recurring?

Triage

SageLogix inc.
Oracle and e-Business Expertise

- Tuning is extremely time-consuming
  - Thorough data-gathering is very time-consuming
  - Don't pay equal attention to every problem
    - There is rarely a *need* to tune everything, anyway… ☺
  - Prioritize!
    - Identify problems that are hurting the most
      - Largest portion of total response time
    - Postpone addressing problems that are not demonstrably serious
      - Find data to match the assertion
  - Don't waste your time and credibility
    - Always work from facts, not hear-say

Triage

SageLogix inc.
Oracle and e-Business Expertise

- Focus on no more than 3-4 problems at a time
  - Understand them completely, fix them, and then re-evaluate!
- For each problem:
  1. Identify
  2. Prioritize
  3. Reproduce/baseline, understand, and try fixes
  4. Implement
  5. Repeat from step 1…
- This method yields dramatic improvements after only a few cycles…

Triage: Identify

SageLogix inc.
Oracle and e-Business Expertise

1. Identify
2. Prioritize
3. Reproduce/baseline, understand, and fix
4. Implement
5. Repeat from step 1…

Two roads diverge…

SageLogix inc.
Oracle and e-Business Expertise

- *Response-time* equals *Service-time* plus *Wait-time*

  R = S + W

  - What else could the RDBMS engine be doing?
    - Anything else is either the hardware, operating-system, I/O subsystem, network, or the application itself
  - Performance problems in the database are either:
    - Resource over-consumption (a.k.a. high *service-time)*
    - Resource contention (a.k.a. high *wait-time)*

Fix or Accommodate?

SageLogix inc.
Oracle and e-Business Expertise

- Fixing the problem:
  - Tuning an inefficient process to make it do less work
  - Find and fix a resource bottleneck
- Accommodating the problem:
  - Add more resources (i.e. bigger machine, more CPUs, RAM disk, faster disk drives)
- Of course it is usually better to fix the problem
  - But sometimes real life doesn't allow this
  - If you accommodate a problem, be sure to understand that it is *not resolution*

So how do I do it?

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- Be familiar with O/S utilities to get an idea about the nature of the problem:
  - Is the problem the database?
  - CONTENTION (i.e. locking, blocking, etc)
    - not very busy, but performance is terrible
  - SATURATION (i.e. resource-hogs)
    - extremely busy, with terrible performance
  - Both? Attend to *saturation* first, then deal with any remaining *contention* later
    - often *saturation* causes *contention*
- Perform YAPP analysis within Oracle
  - Use V$SQLAREA to find *resource-hogs* (saturation), if indicated
  - Use SESSION WAIT views to find *contention*, if indicated

O/S utilities

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- Unix/Linux
  - "top" command
    - to find the heaviest consumers of CPU
    - *top* utility (www.sunfreeware.com), *glance* on HP-UX
    - if you don't have such a utility, you can use:

      ```
      alias pstop="ps -eaf | sort -n +3 | tail"
      ```
  - iostat
    - examine disk I/O utilization using "iostat –x"
  - vmstat, sar
    - examine CPU, memory, swap/page statistics
- Windows NT
  - Task Manager

O/S utilities

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

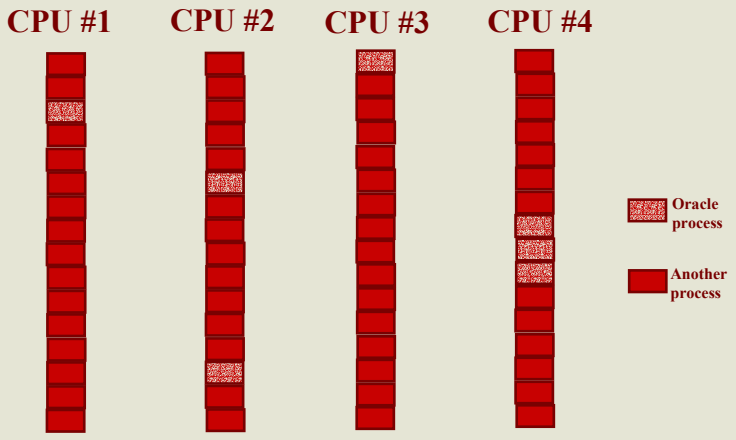

V$ views

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- Oracle records run-time statistics to data structures in memory
  - Available for query via internal X$ tables
  - X$ tables can also be mapped to data structures in Oracle's control files
- V$ views are a more user-friendly view of these statistics
  - Definitions of V$ views can be obtained from V$FIXED_VIEW_DEFINITION view
- Three types of information in V$ views
  - Real-time
  - Cumulative process-, session-, or transaction-based
  - Cumulative instance-based

V$ views

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- Real-time session-based views
  - V$SESSION_WAIT, V$LOCK, V$LATCH_HOLDER, V$SESSION
- Cumulative process-based views
  - V$PROCESS, V$PX_PROCESSES, V$DISPATCHER, V$SHARED_SERVER
- Cumulative session-based views
  - V$SESSION_LONGOPS, V$SESSION_EVENT, V$SESSTAT
- Cumulative transaction-based views
  - V$TRANSACTION

V$ views

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- Cumulative instance-based views
  - V$DATAFILE, V$SYSTEM_EVENT, V$RESOURCE_LIMIT, V$SYSSTAT, V$SGASTAT, V$SQLAREA
  - Just about all of the rest of the views fall into this category…
- Most are purely *memory-based* data structures
  - V$SESSION, V$TRANSACTION, V$SQLAREA, V$LOCK, etc
- Some are based on data structures in control files
  - V$DATAFILE, V$LOGFILE, V$DATABASE, etc

BSTAT/ESTAT

SageLogix inc.
Oracle and e-Business Expertise

- Ancient report based on *snapshots* of information in cumulative instance-based V$ views
  - SQL scripts intended for *Server Manager (svrmgrl)*
    - located in *$ORACLE_HOME/rdbms/admin* directory
  - Execute *utlbstat.sql* to take *beginning snapshot* of views
    - `create table stat$begin_xxx as select * from v$xxx;`
  - Execute *utlestat.sql* to take *ending snapshot* and spool out report (*report.txt*)
    - `create table stat$begin_xxx as select * from v$xxx;`
    - Report is a *delta*
  - Don't try to use BSTAT/ESTAT on a regular basis
    - Capture the time period of a problem only
- Very complete
  - Most important V$ views are queried
  - Extremely detailed
    - *too detailed* for effective analysis

STATSPACK

SageLogix inc.
Oracle and e-Business Expertise

- STATSPACK does essentially the same thing as BSTAT/ESTAT
  - Except HOW it takes *snapshots* of statistics in V$ views
    - Regularly scheduled via DBMS_JOB
    - Stores the information in tables in PERFSTAT schema
  - STATSPACK report contains more information than BSTAT/ESTAT
    - But still difficult to sift through all that information to find a problem

Using STATSPACK

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- Levels of data gathering:
  - Level 0
    - Similar stats as BSTAT/ESTAT (8i+)
      - V$SYSSTAT, V$SYSTEM_EVENT, many more, etc
  - Level 5 (default level)
    - Captures SQL execution info (8i+)
      - Level 0 plus V$SQLAREA info
  - Level 6
    - Captures SQL Plan info (9i+)
      - Level 5 plus V$SQL_PLAN info
  - Level 7
    - Capture segment I/O stats (9i+)
      - Level 6 plus V$SEGMENT_STATISTICS info
  - Level 10
    - Captures latch stats (8i+)
      - Level 7 plus V$LATCH_PARENT, V$LATCH_CHILDREN

Using STATSPACK

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- STATSPACK has only one standard report
  - Named *spreport.sql*
  - Shows *delta info* between two snapshots
    - Just like BSTAT/ESTAT
  - Oracle9*i* also contains an extra detailed report for *drilling down* on the history of a single SQL statement

…that's all folks! So why is STATSPACK so useful?

YAPP Report

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- Web service at http://www.oraperf.com
  - Paper at http://www.oraperf.com/whitepapers.html
  - First became available online 4 years ago
    - Developed by Anjo Kolk
      - Formerly of Oracle RDBMS development (16 yrs), now with Precise
    - Perl-based post-processor web service
    - Makes info from BSTAT/ESTAT and STATSPACK reports easier to analyze
  - Upload *report.txt* or *sp_nnn_nnn.lst* files generated by BSTAT/ESTAT or STATSPACK (respectively)
- MetaLink is rumored to be developing it's own version of a YAPP processor as well

YAPP Report

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- Main body of analysis report is based on the idea of *total response time*
  - *total service time* plus *total wait time*
    - R = S + W
- Analysis report is divided into sections:
  - Response time summary
    - CPU time detailed analysis
    - Wait time detailed analysis
  - "init.ora" parameter values
  - Advise Summary
    - based on info from *response-time* details

YAPP Report

SageLogix inc.

ORACLE AND E-BUSINESS EXPERTISE

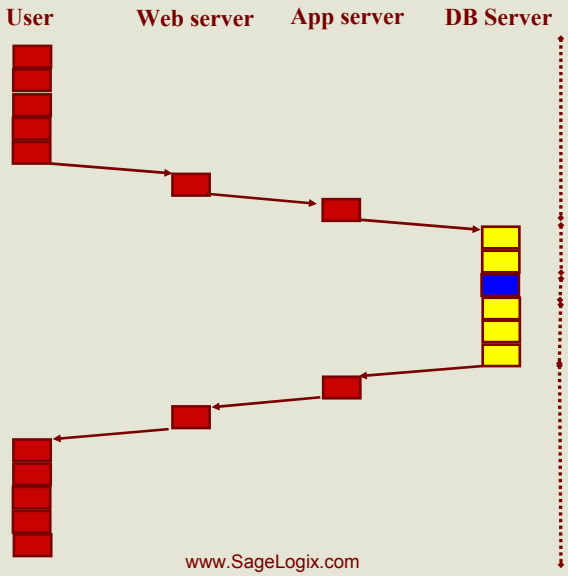

www.SageLogix.com

YAPP Report

SageLogix inc.

ORACLE AND E-BUSINESS EXPERTISE

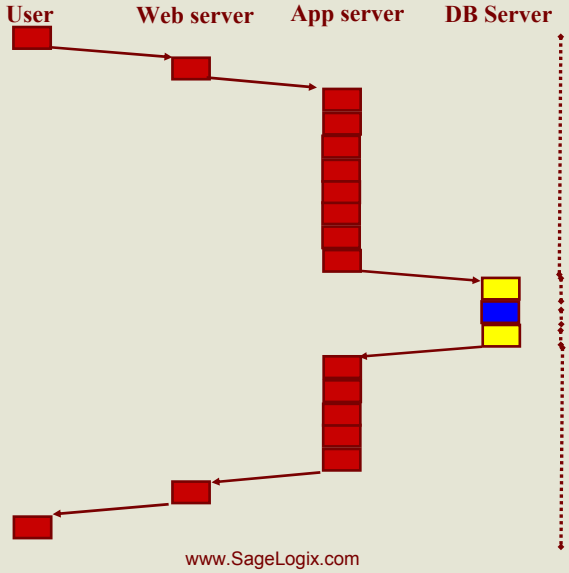

www.SageLogix.com

YAPP Report

SageLogix inc.
Oracle and e-Business Expertise

- Parse SQL
  - High amounts of time spent parsing can be attributed to:
    - Low-level API (OCI, JDBC, precompilers (PRO*C/SQLJ) logic is explicitly parsing for each execute
    - SQL is un-reusable (i.e. bind-variables not being used)
    - Parameter SESSION_CACHED_CURSORS not set
- Recursive SQL
  - High numbers of recursive data dictionary statements (i.e. uncached sequences, SGA too small, etc)
  - PL/SQL packages/procedures/sequences/triggers used but not *pinned*
    - "gen_pin.sql" from http://www.EvDBT.com/tools.htm
  - BEWARE! In some versions of Oracle, all SQL executed within triggers, packages, or stored procedures are counted inside "recursive SQL"

YAPP Report

SageLogix inc.
Oracle and e-Business Expertise

- "Other" SQL processing
  - This should have the majority of time spend processing!
  - Includes SQL *execution* and *fetching*
  - Procedure TOP_STMT2 can be used to detect the *most expensive* SQL statements executed recently
- Wait Events
  - YAPP report provides excellent background information on each *wait event* listed
    - More in the next section…
- "Advise" section at the bottom of the report
  - Summarizes the percentages given for each breakdown to predict potential *maximum gain %* for fixing each problem

Response Time

	Time	Percentage	Per Execute	Per User Call	Per Transactions
Response Time	20985219	100.00%	2.92	4.81	544.72
CPU Time	20367938	97.06%	2.83	4.67	528.69
Wait Time	617281	2.94%	0.09	0.14	16.02

YAPP Report
SageLogix inc.
Oracle and e-Business Expertise

CPU Time

	Time	Percentage	Per Execute	Per User Call	Per Transaction
Total	20367938	100.00%	2.83	4.67	528.69
Parse CPU	46637	0.23%	0.01	0.01	1.21
Recursive CPU	1184109	5.81%	0.16	0.27	30.74
Other CPU	19137192	93.96%	2.66	4.39	496.75

YAPP Report

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

Wait Time

Event	Time	Perc	Per Execute	Per User Call	Per Transaction
latch free	612795	99.27%	0.09	0.14	15.91
enqueue	153	0.02%	0.00	0.00	0.00
log file sync	117	0.02%	0.00	0.00	0.00
buffer deadlock	37	0.01%	0.00	0.00	0.00
write complete waits	5	0.00%	0.00	0.00	0.00
buffer busy waits	5	0.00%	0.00	0.00	0.00

YAPP Report

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

Advise Summary

MaxGain %	What	Detail
23	Reduce the number of buffer gets or executions	Check SQL statement "SELECT /*+ choose */ DS.SHIPMENT_ID, DS.DELIVERY_ID, DT.DEPARTURE_ID FROM DPA_SHIPMENTS DS, DPA_TRUCKS DT WHERE DS.TRUCK_ID = DT" (hash value 73620072).
2	**Tune the cache buffers chain**	No detailed information is available yet
1	**Reduce data block contention**	Check the objects that are inserted into that they have enough freelists

Two roads diverge…

SageLogix inc.
Oracle and e-Business Expertise

- Decision-tree:
  - If *cpu-time* has a larger share of *total-response-time*, then the next step is to look for *inefficient SQL*
  - If *wait-time* has a larger share and I/O represents the majority of waits, then the next step is (again) to look for *inefficient SQL*
  - Otherwise, check the *session-wait interface* for bottlenecks
- First, we will discuss diagnosing service-time
- Then, we will discuss diagnosing wait-time

Shared SQL Area

SageLogix inc.
Oracle and e-Business Expertise

- Shared SQL Area
  - Located in the Library Cache of the Shared Pool in the SGA (system global area)
  - A *cache* of SQL statements (along with their parsed execution plans)
- During the *parse* of a SQL statement
  - the Shared SQL Area must first be checked to see if the SQL statement has already been parsed
    - If entry already exists, it is reused
      - a.k.a. *soft parse*
    - If not, it is parsed and an entry is made in the Shared SQL Area cache
      - a.k.a. *hard parse*

Shared SQL Area

SageLogix inc.
Oracle and e-Business Expertise

- Shared SQL Area is a *hash table* in memory
  - *Hash key* (HASH_VALUE) derived from text of SQL statement
    - Actual algorithm can vary by versions of RDBMS
- Cache maintained by *least-recently-used* (LRU) algorithm
  - When an entry is added, it is added to an "LRU list" (a linked-list) at the *most-recently-used* (MRU) end
  - Each time entry is re-used, it is moved back to the MRU end
  - Lack of reuse gradually slides an entry toward the LRU end, where it may eventually "fall off" and be overwritten

Shared SQL Area

SageLogix inc.
Oracle and e-Business Expertise

- Memory structure is exposed via V$SQLAREA view:
  - ADDRESS
  - HASH_VALUE
  - SQL_TEXT (1st 1000 characters only)
  - DISK_READS
  - BUFFER_GETS
  - ROWS_PROCESSED
  - EXECUTIONS
  - LOADS
  - FIRST_LOAD_TIME
  - COMMAND_TYPE
  - OPTIMIZER_MODE
  - …more…

Shared SQL Area

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- Related views:
  - V$SQL
    - Detail of V$SQLAREA
    - V$SQLAREA is a summary of V$SQL
  - V$SQLTEXT
    - Contains full text of SQL statement
      - Each row has a 64-byte *piece* of the text
  - V$SQLTEXT_WITH_NEWLINES
  - V$SQL_SHARED_MEMORY
    - Detail related to each entry in V$SQLAREA
    - Physical shared memory info

Shared SQL Area

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- statements consuming lots of resources, look in V$SQLAREA for:
  - disk_reads (a.k.a. *physical reads*)
  - buffer_gets (a.k.a. *logical reads*)
  - derived "load" factor for sorting

    ```
    ((disk_reads * 50) + buffer_gets)
    ```
- ideal for creating a top 10 list

```
SELECT    SQL_TEXT, DISK_READS, BUFFER_GETS
FROM      V$SQLAREA
WHERE     DISK_READS > ?????
AND       BUFFER_GETS > ??????
ORDER BY "load factor" formula (above) DESC
```

Shared SQL Area

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- Procedure TOP_STMT is based on a similar statement
  - Available from http://www.EvDBT.com/tools.htm
  - Produces “top 10 offensive SQL” report
- Procedure TOP_STMT2 is very similar, but *aggregates* SQL statements based on the first 60 characters of text
  - Attempts to aggregate statistics from SQL statements which are alike for the first 60 characters
    - Avoiding differences in data values embedded in the WHERE clause

TOP_STMT2 report

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

```
Date:                          03/26/01 10:39:20               Page    2
Database startup:              03/18/01 04:06:03
Total Logical Reads:              91,988,081,578
Total Physical Reads:              1,894,402,241   ("Hit Ratio":  97.94%)
.
SQL Statement Text
------------------
0     SELECT COUNT(1)   FROM RA_SITE_USES SU,RA_ADDRESSES ADDR  WHERE
1     SU.ADDRESS_ID = ADDR.ADDRESS_ID  AND SU.SITE_USE_CODE = :b1  AND
2      SU.LOCATION = :b2  AND ADDR.CUSTOMER_ID = :b3
:         Disk             Buffer                                    Load
:         Reads              Gets     Sorts      Runs     Loads      Factor
:         -----            ------     -----      ----     -----      ------
:    19,235,856       426,258,235         1    29,104         5   2349843.8
:      (1.015%)          (0.463%)
```

Two roads diverge…

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- A look at a sample real-life YAPP report…
- Next, we discuss diagnosing wait-time…

Session Waits

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

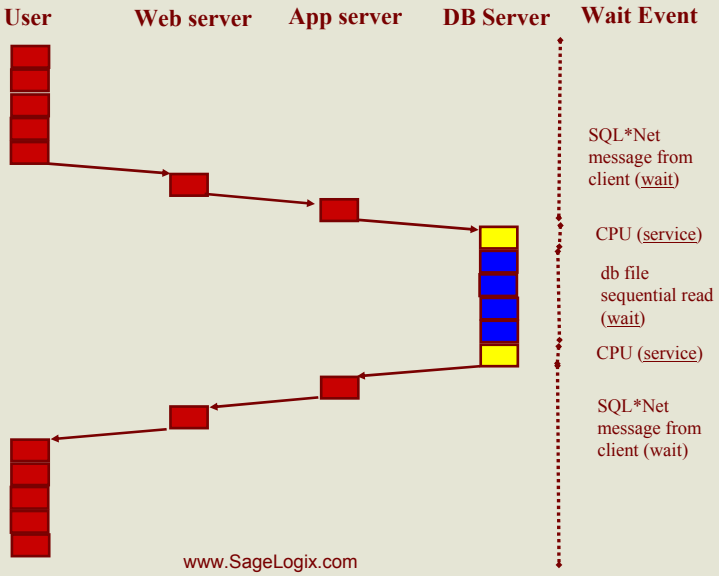

Session Waits

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

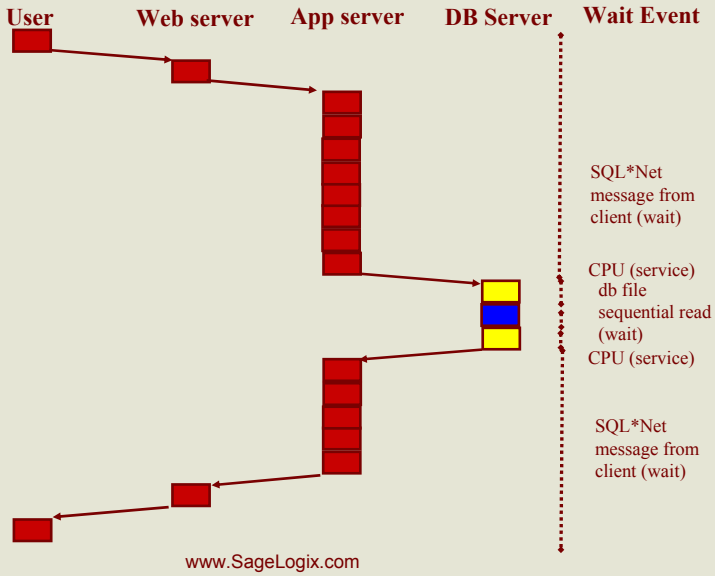

Session Waits

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- Added in Oracle7 v7.0.12
  - Whenever the Oracle RDBMS is going to wait on something for which the session will *lose the CPU*
    - I/O call, lock/sleep, latch/sleep, etc
  - Each time something is about to be called:
    1. Post an event
       - If TIMED_STATISTICS is TRUE, then record current timestamp
    2. Make the call
    3. Upon return from the call:
       - If TIMED_STATISTICS is TRUE, then determine determine elapsed time
    4. Increment *counters* in V$ views
       - If TIMED_STATISTICS is TRUE, then add elapsed time as well

Session Waits

SageLogix inc.
Oracle and e-Business Expertise

- Just four V$ views
  - Three with different levels of *data accumulation*
    - V$SESSION_WAIT
      - Real-time
      - If a session is waiting, then 1 row, else no row for that session
    - V$SESSION_EVENT
      - Cumulative wait-event statistics for all active sessions
      - One row per session per wait-event
    - V$SYSTEM_EVENT
      - Cumulative wait-event statistics for instance
      - One row per wait-event
  - One for *online documentation*
    - V$EVENT_NAME

V$SESSION_WAIT

SageLogix inc.
Oracle and e-Business Expertise

- Real-time view of the SESSION WAIT interface
  - SID
    - PK; FK to V$SESSION view (i.e. client-process information)
  - EVENT
    - event that the session is waiting on
      - 150+ wait-events in Oracle8
      - 217 wait-events in Oracle8*i* R3
      - 361 wait-events in Oracle9*i* R2
      - listed in V$EVENT_NAME view
  - P1TEXT, P1
  - P2TEXT, P2
  - P3TEXT, P3
    - descriptive text and values to further describe event

V$SESSION_WAIT

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- STATE
  - WAITING (session currently waiting)
  - WAITED UNKNOWN TIME (timed_statistics = FALSE)
  - WAITED SHORT TIME (too short - unable to measure)
  - WAITED KNOWN TIME (wait finished)
- SECONDS_IN_WAIT
  - Duration of wait in seconds when STATE = "WAITING"
    - otherwise ignore
- WAIT_TIME
  - Duration of wait in seconds when STATE = "WAITED KNOWN TIME"
    - otherwise ignore

V$SESSION_EVENT

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- V$SESSION_EVENT has cumulative totals of *wait event info* for connected sessions
  - SID (foreign key to V$SESSION view)
  - EVENT (foreign key to V$EVENT_NAME view)
  - TOTAL_WAITS
  - TOTAL_TIMEOUTS
  - TOTAL_TIME
  - AVERAGE_TIME
  - MAX_WAIT
    - Last 3 columns not populated if parameter TIMED_STATISTICS set to FALSE

V$SYSTEM_EVENT

SageLogix inc.
Oracle and e-Business Expertise

- V$SYSTEM_EVENT has cumulative totals of *wait event info* for entire database instance since startup
  - EVENT
  - TOTAL_WAITS
  - TOTAL_TIMEOUTS
  - TOTAL_TIME
  - AVERAGE_TIME
    - Last 2 columns not populated if parameter TIMED_STATISTICS set to FALSE

V$EVENT_NAME

SageLogix inc.
Oracle and e-Business Expertise

- EVENT#
  - Never used in other views
- NAME
  - Name used in other views
- PARAMETER1
- PARAMETER2
- PARAMETER3
  - Terse descriptive text, same values as P1TEXT, P2TEXT, P3TEXT in V$SESSION_WAIT

Using Session Wait views

SageLogix inc.
Oracle and e-Business Expertise

- Query V$SESSION_WAIT get a detailed idea of what is happening *right now*, right this *second* for active sessions
- Query V$SESSION_EVENT to get an idea of what the currently active sessions have been doing
  - Either right now or the *recent past*
- Query V$SYSTEM_EVENT to get a high-level overview of what the entire instance has been bottlenecking upon

Common wait events

SageLogix inc.
Oracle and e-Business Expertise

db file scattered read	I/O request usually associated with FULL table scans • Multi-block sequential reads
db file sequential read	I/O request usually associated with indexed table scans • Single-block random reads
latch free	Waiting for a latch sleep to complete
free buffer waits	Waiting on DBWR to clear *dirtied* blocks
write complete waits	Waiting on DBWR to finish writing a buffer to disk
buffer busy waits	Waiting on another process which has a buffer locked for update
log file sync	Waiting on LGWR to finish writing a redo batch from Log Buffer to disk
enqueue	Waiting upon an application lock a.k.a. enqueue

Query v$session_wait

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

```
SQL> col event format a20 truncate
SQL> col p1 format a12 truncate
SQL> col p2 format a12 truncate
SQL> col p3 format a12 truncate
SQL>
SQL> select event,
  2         p1text || '=' || p1 p1,
  3         p2text || '=' || p2 p2,
  4         p3text || '=' || p3 p3
  5  from v$session_wait
  6  where event not like '% timer'
  7  and event not like 'rdbms ipc message';
```


Query v$session_wait

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

```
  SID   EVENT                P1           P2           P3
------ -------------------- ------------ ------------ ---------
   127 latch free           addr=4597239 latch=30     =0
   201 db file scattered re file=73      block=75444  blocks=16
   352 db file sequential r file=111     block=9291   blocks=1
```

- Become familiar with the contents of *Appendix A* of the *Oracle8 Server Reference* manual to understand what these events could mean
- Thorough understanding of Oracle process architecture, enqueues, latches, and the Shared Pool is useful

Useful scripts

SageLogix inc.
Oracle and e-Business Expertise

- On http://www.EvDBT.com/tools.htm:
  - latchfree.sql
    - Queries V$SESSION_WAIT for “latch free” event, joins to V$SESSION, V$SQLAREA, and V$LATCH_CHILDREN
  - sw.sql
    - Queries V$SESSION_WAIT for specified event name, also displays V$SESSION_EVENT info for comparison
  - sysevent.sql
    - Queries V$SYSTEM_EVENT, sorts by TIME_WAITED, estimates “percentage of concern”
  - w.sql
    - Quick summary of V$SESSION_WAIT information
  - wts.sql
    - Queries V$SESSION_WAIT for “buffer busy waits”, “free buffer waits”, and “write complete waits”. Joins to DBA_EXTENTS to find segment info, V$SQLTEXT for SQL, V$SESSION

sysevent.sql report

SageLogix inc.
Oracle and e-Business Expertise

Event Name	Total Waits (in 1000s)	Total Timeouts (in 1000s)	Time Waited (in Hours)	% of Concern	Avg Wait (Secs)
enqueue	0.99	0.99	0.83	53.40	3.01
control file parallel writ	108.73	0.00	0.36	23.33	0.01
direct path write	62.61	0.00	0.23	14.56	0.01
direct path read	31.17	0.00	0.06	3.92	0.01
log file parallel write	2.94	0.00	0.02	1.55	0.03
db file parallel write	0.83	0.00	0.01	0.96	0.07
db file sequential read	52.87	0.00	0.01	0.90	0.00
log file sync	2.02	0.00	0.01	0.80	0.02
SQL*Net message from clien	247.38	0.00	679.24	0.00	9.88
rdbms ipc message	335.40	329.24	559.64	0.00	6.01
pmon timer	109.24	109.24	93.37	0.00	3.08
smon timer	1.10	1.09	93.35	0.00	305.78

w.sql report

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

Event	#Sess
enqueue	12
db file scattered read	6
latch free	1
db file sequential read	1
buffer busy waits	1

sw.sql report

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

```
Enter value for event: enqueue

                    Current                                    Sum        Sum        Sum
     Current          Event Current  Summary                 Time        Avg        Max
 Sid Event             Seq# State    Events                  Waited     Waited     Waited
---- --------------- ------- -------- --------------- -------- --------- --------
 210 enqueue name|mo      62 WAITING  SQL*Net              2133      10.78       17
     mode=TX id1=105         (168619 message from
     033 id2=0               s)      client
     1413697536,
     #bytes = 1

                                     SQL*Net              1897      63.33       64
                                     message to
                                     client

                                     SQL*Net                 0       0.00        0
                                     break/reset to
                                     client
```

www.SageLogix.com

latchfree.sql report

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

Sid	Latch	Program	Clnt PID	Srvr PID	SQL Text
210	cache buffers chains (child# 375)	sqlplus (TNS V1-V3) @rpt_xref5	curley 20777	1863	select v1.VALUE_STRING as "DeaNumber",v2.VALUE_STRING as "LastName" from ATTRIBUTE_VALUES v1, ATTRIBUTE_VALUES v2 where v1.ATTRIBUTE_DEFINITION_ID = 16 and v1.VALUE_STRING = 'A90527638' and v1.SOURCE_CODE = v2.SOURCE_CODE and v1.SOURCE_CODE = 'TXR' and v2.ATTRIBUTE_DEFINITION_ID = 6
655	cache buffers chains (child# 2 2857)	sqlldr (TNS V1-V3)	larry 11896	9973	insert into ATTRIBUTE_VALUES values (:b1, :b2, :b3, :b4, :b5, :b6, :b7, :b8, :b9,

Further advice

SageLogix inc.
ORACLE AND E-BUSINESS EXPERTISE

- Event “latch free”:
  - Use scripts “latches.sql” and “latchfree.sql” for more info…
    - “Cache buffers chains” latch is most commonly contended
    - “library cache” and “shared pool” high usually indicates parse problems (i.e. bind-variable issue)
- Event “enqueue”:
  - Use scripts “locks.sql”, “sess_locks.sql”, and “enqwaits.sql” for more info
  - Enqueues leave *no* history in Oracle, so you have to *catch them in the act*
- Events “free buffer waits”, “buffer busy waits”, and “write complete waits”
  - Use “wts.sql”
  - Be aware of possible I/O problems with DBWR

www.SageLogix.com

Advice for DBAs

SageLogix inc.
Oracle and e-Business Expertise

- Be sure to set TIMED_STATISTICS to TRUE
  - Otherwise, the timing information in all of the *session wait* views will not be available

    …and the information is pretty much *useless* without timing information!!!
  - Some people think that they are *optimizing* performance by removing what scant *processing overhead* is implied by using TIMED_STATISTICS
    - This is a case of *"penny wise, pound foolish"*
    - *How can you tune something if you can't gather useful information?*

More useful scripts

SageLogix inc.
Oracle and e-Business Expertise

http://www.ixora.com.au/scripts/waits.htm:

- response_time_breakdown.sql
- session_times.sql
- resource_waiters.sql
- trace_waits.sql
- waiters.sql
- resource_waits.sql
- routine_waits.sql
- system_times.sql

Useful Documentation

SageLogix inc.
Oracle and e-Business Expertise

- Craig Shallahamer's "*Direct Contention Identification using Oracle's Session Wait Tables*"
  - Available at http://www.OraPub.com - "technical papers" link
- Craig Shallahamer's "*Oracle Performance Triage: Stop the Bleeding*"
- Mogens Norgaard "*Introducing Oracle's Wait Interface"*
  - Available at http://www.hotsos.com/catalog/level.html

Online resources

SageLogix inc.
Oracle and e-Business Expertise

- Oak Table http://www.oaktable.net/
  - IXORA - Steve Adams http://www.ixora.com.au/
  - Hotsos - Cary Millsap http://www.hotsos.com/
  - Jonathan Lewis http://www.jlcomp.demon.co.uk/
  - Anjo Kolk http://www.oraperf.com/
  - Miracle A/S - Mogens Norgaard http://www.miracleas.dk/
  - SageLogix - Tim Gorman http://www.evdbt.com/ and http://www.sagelogix.com
- Orapub - Craig Shallahamer http://www.orapub.com/
- Oracle FAQ http://www.orafaq.com/
  - Subscribe instructions and archive for ORACLE-L list
- UB Tools - Danisment Gazi Unal http://www.ubtools.com/

Triage: Prioritize

SageLogix inc.
Oracle and e-Business Expertise

1. Identify
2. Prioritize
3. Reproduce/baseline, understand, and fix
4. Implement
5. Repeat from step 1…

Prioritize

SageLogix inc.
Oracle and e-Business Expertise

- YAPP report provides guidance
  - Also helps *quantify* the expected benefit from fixing an identified problem
- Must match identified problems with important business processes
  - No kudos for fixing a problem that nobody is concerned about!

Reproduce/baseline

SageLogix inc.
Oracle and e-Business Expertise

- Reproduce and baseline the problem
  - Utilize SQL Trace level-8 on a test process
    - To gather both service-time and wait-time information from the specific process with the problem
- Understand and fix the problem
  - Once you have a reproducible test case, you can try to fix it
    - With the original test case and the fix, you can demonstrate quantifiable benefits

Implement

SageLogix inc.
Oracle and e-Business Expertise

- Deploying the fix in production
  - Possibly the hardest step of all
    - Rules of change management vary
      - Some shops are very strict
      - Some shops are not
    - Knowing what the fix is doesn't matter if you cannot deploy it
  - Having *quantified* the benefit of the fix will open many doors

Summary

SageLogix inc.
Oracle and e-Business Expertise

- Think *forensically*
  - Most *eyewitnesses* are extremely unreliable
    - Polite skepticism, not annoying cynicism
  - Most *experts* are *not*
    - If they don't have a grasp of the facts for this situation
  - Allow the *empirical facts* to speak
    - Use testimony from *eyewitnesses* and *experts* to refine the story that the *facts* tell
    - Approach the problem from a high-level and quickly identify and prioritize the major issues
      - Drill down on a *very small number* of the *very worst issues*

Q&A

SageLogix inc.
Oracle and e-Business Expertise