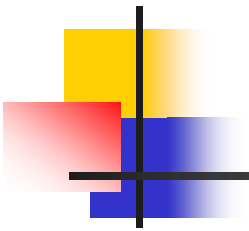

Oracle Wait Interface: What, Why and How

Kirtikumar Deshpande

NYOUG

June 6, 2006

About Me

- Senior Oracle DBA
  - Verizon Information Services
    - Phone Directories Publication

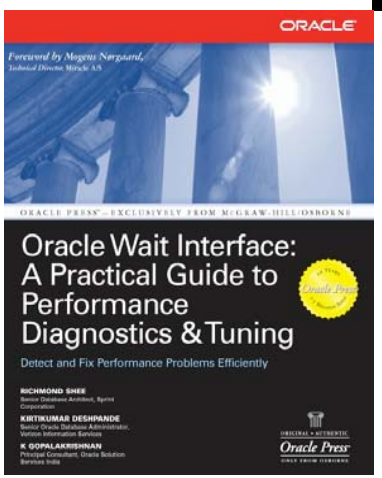

Agenda

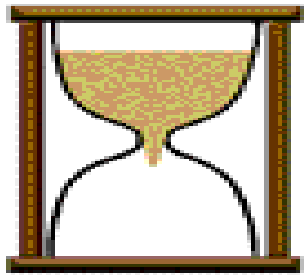

- What is OWI
- Why OWI
- How to use OWI
- Q & A

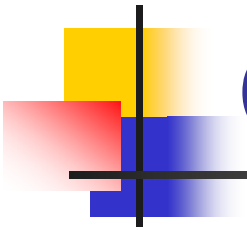

Oracle Wait Interface

What is it?

A tool set that tracks wait events and time waited, to identify bottlenecks throughout the life of a session.

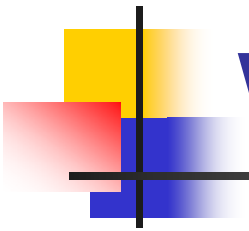

What is a Wait Event?

- Named section of Oracle kernel function
- Processes wait for
  - Resource availability
  - Completion of an action
  - Work to do
- OWI enables measurement of such waits

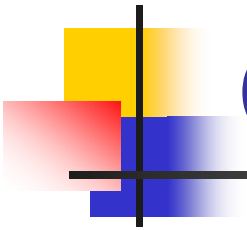

Oracle Wait Interface

Why OWI?

Specific bottleneck areas can quickly be identified and improvements can be targeted to those areas.

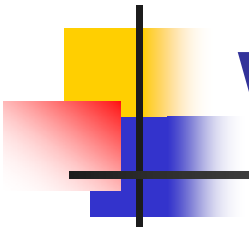

Why OWI

Response Time = Service Time + Wait Time

- Makes logical sense to use response time
- Provides a Methodical Approach to Performance Troubleshooting, Tuning & Monitoring
- Pinpoints the exact location of the bottlenecks
- Possibly helps improve Service Time
- Helps improve Response Time

File Edit View Favorites Tools Help Address http://dfw2sn327:5500/em/console/database/instance/sitemap?event=doLoad&target=BVWF Go Back

ORACLE Enterprise Manager 10g
Database Control

Setup Preferences Help Logout

Database

Logged in As SYSTEM

Database: BVWP.WORLD

Home | **Performance** | Administration | Maintenance

Click on an area of a graph or legend to get more detail.

View Data: Real Time: 15 Second Refresh

Best viewed using latest SVG plugin

Host

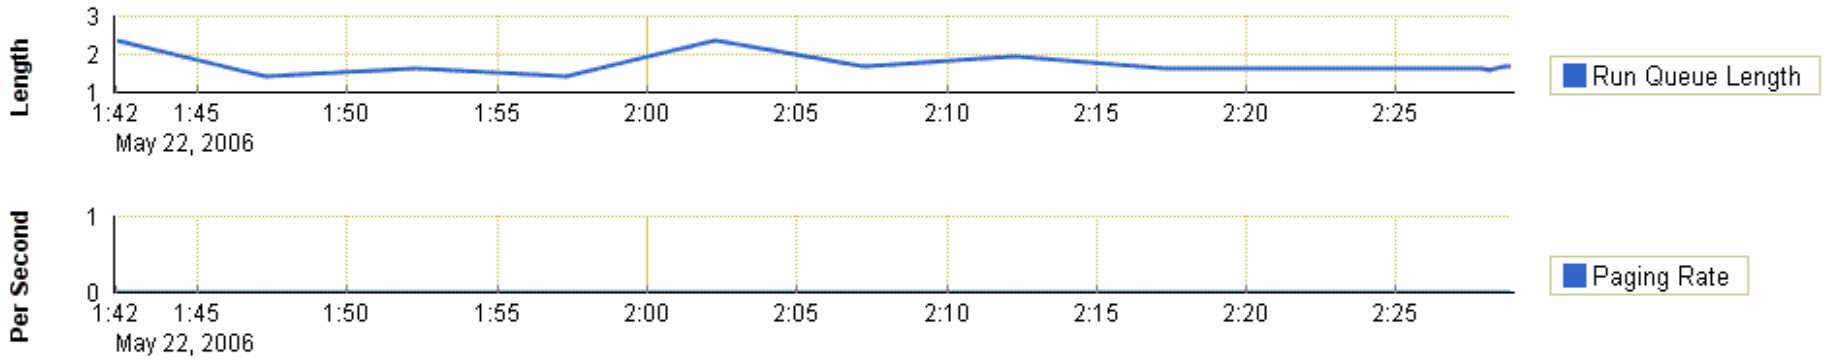

Sessions: Waiting and Working

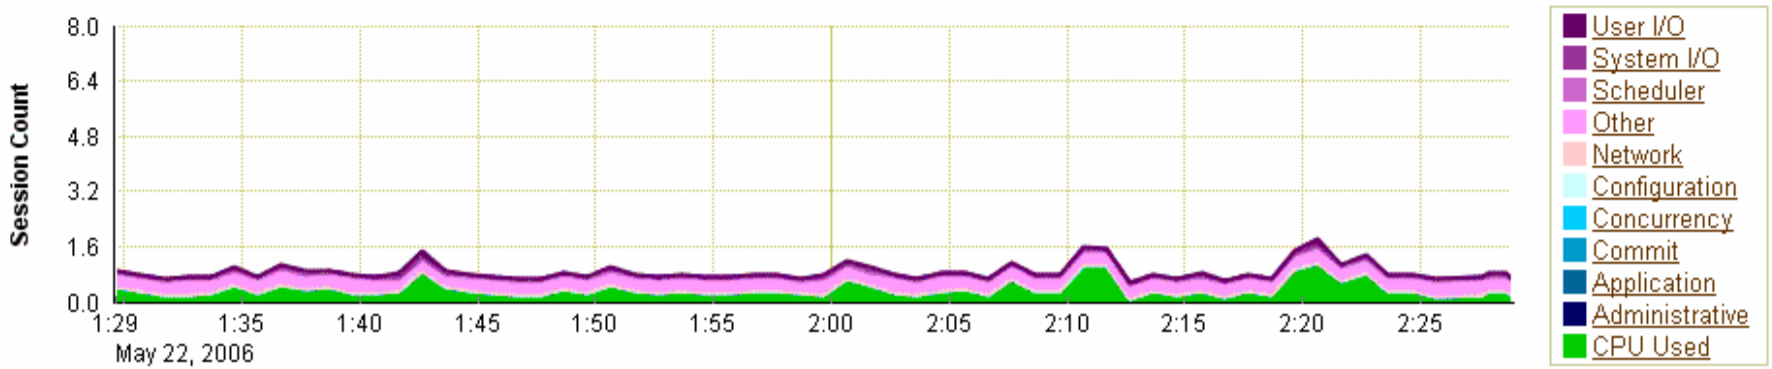

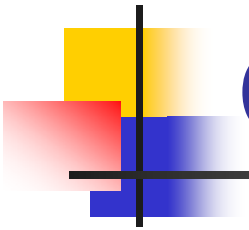

OWI

- Introduced in Oracle 7.0.12
- Set of Views (foundation)
  - V$SYSTEM_EVENT
  - V$SESSION_EVENT
  - V$SESSION_WAIT
  - V$EVENT_NAME
- File generated by tracing session with event 10046 (Extended SQL Trace)

Number of Wait Events

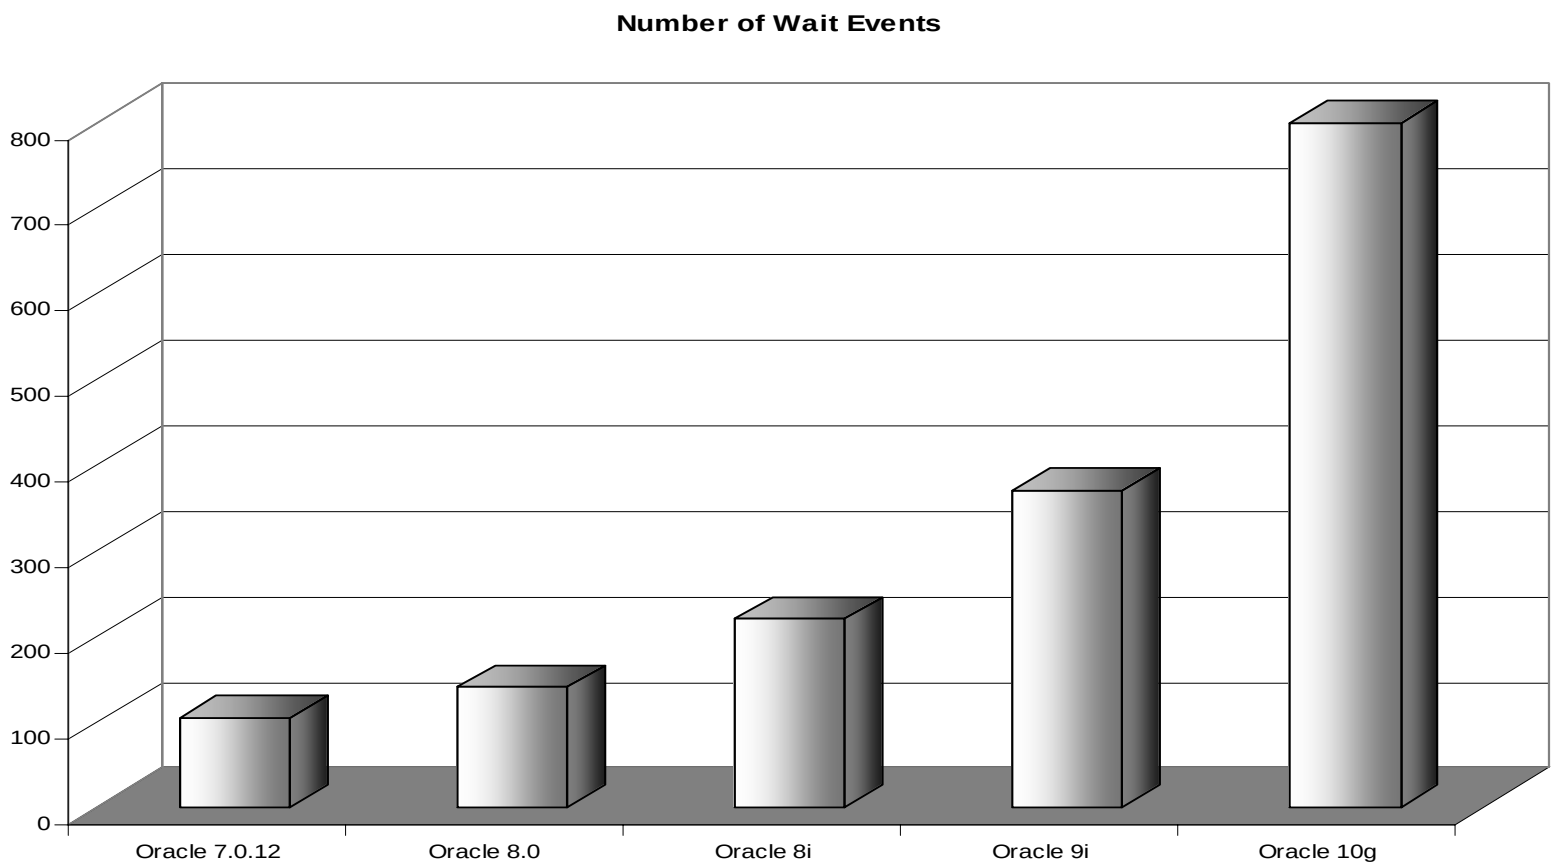

OWI Components

Up to Oracle9i R2

V$EVENT_NAME

V$SYSTEM_EVENT

V$SESSION_EVENT

V$SESSION_WAIT

Event 10046 TRACE

V$STATNAME

V$SESSTAT

TIMED_STATISTICS=TRUE

Oracle10g Release 1

V$EVENTMETRIC

V$EVENT_HISTOGRAM

V$SERVICE_EVENT

V$ACTIVE_SESSION_HISTORY

V$SYSTEM_WAIT_CLASS

V$SESSION_WAIT_CLASS

V$SERVICE_WAIT_CLASS

V$SESSION_WAIT_HISTORY

V$WAITCLASSMETRIC

V$WAITCLASSMETRIC_HISTORY

V$SYSMETRIC_HISTORY

V$SERVICEMETRIC_HISTORY

DBA_HIST_EVENT_NAME

DBA_HIST_SYSTEM_EVENT

DBA_HIST_BG_EVENT_SUMMARY

DBA_HIST_WAITCLASSMET_SUMMARY

DBA_HIST_SERVICE_WAIT_CLASS

DBA_HIST_ACTIVE_SESS_HISTORY

OWI Components

V$EVENT_NAME
Events Defined in the Database

Name	Type
EVENT#	NUMBER
EVENT_ID	NUMBER
NAME	VARCHAR2(64)
PARAMETER1	VARCHAR2(64)
PARAMETER2	VARCHAR2(64)
PARAMETER3	VARCHAR2(64)
WAIT_CLASS_ID	NUMBER
WAIT_CLASS#	NUMBER
WAIT_CLASS	VARCHAR2(64)

(Red - From Oracle 10g)

V$SESSION_WAIT
Currently Waiting Sessions

Name	Type
SID	NUMBER
SEQ#	NUMBER
EVENT	VARCHAR2(64)
P1TEXT	VARCHAR2(64)
P1	NUMBER
P1RAW	RAW(8)
P2TEXT	VARCHAR2(64)
P2	NUMBER
P2RAW	RAW(8)
P3TEXT	VARCHAR2(64)
P3	NUMBER
P3RAW	RAW(8)
WAIT_TIME	NUMBER
SECONDS_IN_WAIT	NUMBER
STATE	VARCHAR2(19)

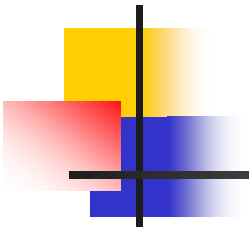

STATE in V$SESSION_WAIT

- WAITING
  - Session is currently waiting. SECONDS_IN_WAIT shows wait time.
- WAITED SHORT TIME
  - Session waited for insignificant amount of time before acquiring required resource. Less than a centi-second. WAIT_TIME = –1.
- WAITED KNOWN TIME
  - Session waited for WAIT_TIME before acquiring required resource.
- WAITED UNKNOWN TIME
  - TIMED_STATISTICS is not set to TRUE. WAIT_TIME = –2.

V$SESSION_EVENT
Waited Events for Current Sessions

Name	Type	
SID	NUMBER	
EVENT	VARCHAR2(64)	
TOTAL_WAITS	NUMBER	
TOTAL_TIMEOUTS	NUMBER	
TIME_WAITED	NUMBER	
AVERAGE_WAIT	NUMBER	
MAX_WAIT	NUMBER	
TIME_WAITED_MICRO	NUMBER	
EVENT_ID	NUMBER	(10g)

V$SYSTEM_EVENT
Waited Events for the Instance

Name	Type	
EVENT	VARCHAR2(64)	
TOTAL_WAITS	NUMBER	
TOTAL_TIMEOUTS	NUMBER	
TIME_WAITED	NUMBER	
AVERAGE_WAIT	NUMBER	
TIME_WAITED_MICRO	NUMBER	
EVENT_ID	NUMBER	(10g)

V$SESSION_WAIT_HISTORY (10g)

Shows 10 most recent waits for session

SID	NUMBER
SEQ#	NUMBER
EVENT#	NUMBER
EVENT	VARCHAR2(64)
P1TEXT	VARCHAR2(64)
P1	NUMBER
P2TEXT	VARCHAR2(64)
P2	NUMBER
P3TEXT	VARCHAR2(64)
P3	NUMBER
WAIT_TIME	NUMBER
WAIT_COUNT	NUMBER

V$SESSION_WAIT_HISTORY (10g)
Example

SID	SEQ#	EVENT	P1	P2	P3	W_TM	W_COUNT
40	1	log file sync	962	0	0	0	1
40	2	SQL*Net message from client	1650815232	1	0	8721	1
40	3	SQL*Net message to client	1650815232	1	0	0	1
40	4	db file scattered read	5	225	4	0	1
40	5	db file scattered read	5	217	8	0	1
40	6	db file scattered read	5	209	8	0	1
40	7	db file scattered read	5	201	8	0	1
40	8	db file scattered read	5	193	8	1	1
40	9	db file scattered read	5	185	8	0	1
40	10	SQL*Net message to client	1650815232	1	0	0	1

V$EVENT_HISTOGRAM (10g)

Histogram of # of waits, maximum waits, and wait time

EVENT#	EVENT	WAIT_TIME_MILLI	WAIT_COUNT
294	db file sequential read	1	29453
294	db file sequential read	2	7319
294	db file sequential read	4	749
294	db file sequential read	8	904
294	db file sequential read	16	3395
294	db file sequential read	32	2167
294	db file sequential read	64	270
294	db file sequential read	128	56
294	db file sequential read	256	16
294	db file sequential read	512	9
295	db file scattered read	1	13397
295	db file scattered read	2	2202
295	db file scattered read	4	881
295	db file scattered read	8	822
295	db file scattered read	16	672
295	db file scattered read	32	523
295	db file scattered read	64	228

V$SYSTEM_WAIT_CLASS (10g)

Shows instance-level waits by class

WAIT_CLASS_ID	WC#	WAIT_CLASS	TOTAL_WAITS	TIME_WAITED
3875070507	4	Concurrency	270	1596
3290255840	2	Configuration	378	1222
4217450380	1	Application	3680	321401
1893977003	0	Other	10394	65063
3386400367	5	Commit	28164	17657
1740759767	8	User I/O	72340	22764
2000153315	7	Network	79141	200
4108307767	9	System I/O	206748	174050
2723168908	6	Idle	1156193	364114196

Oracle10g Wait Classes

CLASS	NBR_EVENTS
Administrative	42
Application	10
Cluster	45
Commit	1
Concurrency	17
Configuration	21
Idle	58
Network	25
Other	556
Scheduler	2
System I/O	19
User I/O	12

Oracle10g Wait Classes

- Administrative (DBA tasks affecting others)
  - Buffer pool resize, Offline rollback segment, Index rebuilds
- Application
  - Row level locks or explicit locks
- Commit
  - Redo log write confirmation
- Concurrency
  - Buffer busy waits, Library cache latch
- Configuration
  - Sizing for log buffer, log files, SGA, Contention for ST enqueue
- Network
  - More data to/from database link/client, Remote archive destination
- User I/O
  - Database file reads, single writes, Direct path reads/writes, BFILE reads
- System I/O
  - Redo log writes, Archiving redo logs, Control file writes
- Idle
  - Indicates that session is inactive, waiting for more work

Event 10046 TRACE

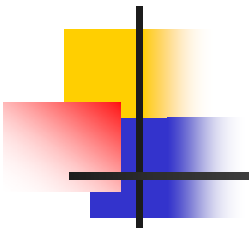

- Similar to setting SQL_TRACE = true;
- Referred to as Extended SQL Trace
- Trace Level setting controls the information reported
- Generated trace file can be processed using *tkprof*
- TIMED_STATISTICS = TRUE
- MAX_DUMP_FILE_SIZE = unlimited

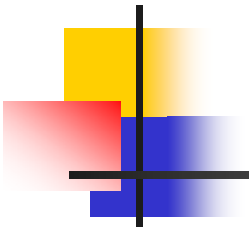

Event 10046 Trace

- Trace levels:
  - Level 1 = Enable standard sql_trace functionality (default)
  - Level 4 = Level 1 plus bind variable values
  - Level 8 = Level 1 plus wait events
  - Level 12 = Level 1 plus bind variables values plus wait events

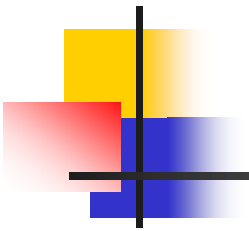

Event 10046 Trace

- Trace your own Session:
  - *alter session set events '10046 trace name context forever, level 8';*
  - Run your SQL script or program ....
  - *alter session set events '10046 trace name context off';*
  - Look for the trace file in UDUMP directory
  - *alter session set tracefile_identifier = 'mytrace';*

Event 10046 Trace

- Trace your own Session:
  - exec sys.dbms_support.start_trace;
    -- By default Wait Event information is written to the trace file.
  - exec sys.dbms_support.start_trace(waits => TRUE, binds=> TRUE);
  - Run your SQL script or program
  - exec sys.dbms_support.stop_trace;

(You must run *dbmssupp.sql* to install *dbms_support* package)

Event 10046 Trace

- Trace other Session:
  - exec dbms_support.start_trace_in_session(
    sid => 1234,
    serial# => 56789,
    waits => true,
    binds => true);
  - Run SQL script or program in other session, if not already running
  - exec dbms_support.stop_trace_in_session(
    sid => 1234,
    serial# => 56789);

Event 10046 Trace (10g)

- Trace other Session:
  - exec dbms_monitor.session_trace_enable(
    session_id => 1234,
    serial_num => 56789,
    waits => true,
    binds => true);
  - Let the session execute SQL script or program for some amount of time
  - exec dbms_monitor.session_trace_disable(
    session_id => 1234,
    serial_num => 56789);

Event 10046 Trace

- What's in the Trace File?

```
PARSING IN CURSOR #1 len=923 dep=0 uid=82 oct=3 lid=82 tim=1071461386936456 hv=3471484162 ad='db203a8'
select y.oppar_db_job_name ,y.oppar_db_job_rec ,y.oppar_db_prefix ,y.oppar_db_request_flag ,y.oppar_db_run_id
      ,TO_CHAR(y.oppar_db_last_date,'yyyymmdd') ,oppar_run_mode from
     ..................................
END OF STMT
EXEC#1:c=2720000,e=2819768,p=29022,cr=31542,cu=0,mis=0,r=0,dep=0,og=4,tim=1071461386936431
FETCH #1:c=0,e=9,p=0,cr=0,cu=0,mis=0,r=0,dep=0,og=4,tim=1071461386936555
WAIT #1: nam='SQL*Net message to client' ela= 5 p1=1952673792 p2=1 p3=0
*** 2004-10-07 14:14:40.246
WAIT #1: nam='SQL*Net message from client' ela= 19535208 p1=1952673792 p2=1 p3=0
BINDS #1:
 bind 0: dty=1 mxl=32(03) mal=00 scl=00 pre=00 oacflg=00 oacfl2=1 size=32 offset=0
   bfp=110319ed0 bln=32 avl=00 flg=05
WAIT #1: nam='db file sequential read' ela= 27 p1=45 p2=119835 p3=1
WAIT #1: nam='db file sequential read' ela= 10 p1=45 p2=119838 p3=1
WAIT #1: nam='db file sequential read' ela= 11 p1=45 p2=119841 p3=1
WAIT #1: nam='db file sequential read' ela= 10 p1=45 p2=119843 p3=1
WAIT #1: nam='db file scattered read' ela= 74 p1=45 p2=119847 p3=2
WAIT #1: nam='db file sequential read' ela= 9 p1=45 p2=119852 p3=1
```

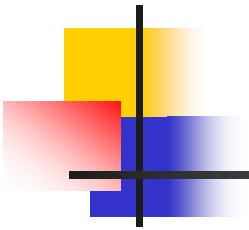

Most Common Wait Events

- db file sequential read
  - Usually Index read - 1 Oracle block
- db file scattered read
  - Full table scan, Fast full index scan - multiple Oracle blocks (db_file_multiblock_read_count)
- buffer busy waits
  - Multiple sessions requesting the same block that has to be read from disk.
  - Multiple sessions waiting for a change to complete in the same block.

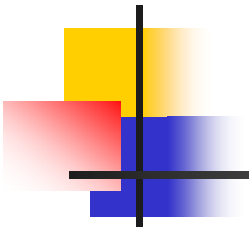

Most Common Wait Events

- free buffer waits
  - DBWR not keeping up with free buffer demands
  - Buffer cache is too small
- log file sync
  - Transaction termination (commit) - the foreground process is waiting for LGWR
  - Large log buffer
- enqueue
  - Wait for a local lock
  - Contention for a specific row in the table – TX lock
  - Contention for the ST lock
  - Wait for INITRANS
  - 184 individual enqueue wait events in Oracle10g Release 1

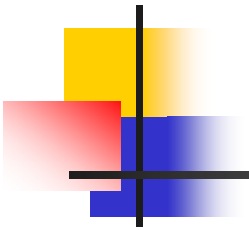

Most Common Wait Events

- latch Free
  - Competition for objects that are protected by serialized mechanism
  - 28 individual latch wait events in Oracle10g Release 1
- SQL*Net break/reset to client/dblink
- SQL*Net more data to client/dblink
- SQL*Net message to client/dblink
- SQL*Net more data from client/dblink
  - Network latency
  - Bottlenecks in client program
  - Bottlenecks in remote server

Q & A

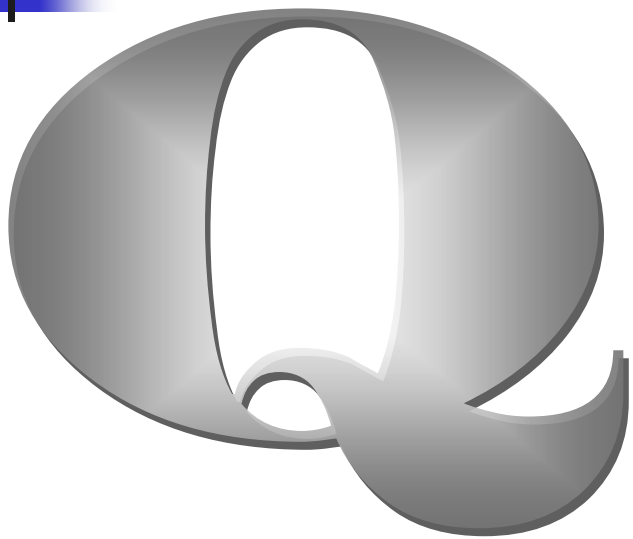

Kirtikumar_Deshpande@yahoo.com