# Optimizing Oracle 12c Performance

## Full Database Caching and Database In-Memory Option

TidalScale™

Sharding Motherboards Since 2013

# daniel.morgan@tidalscale.com

- Director of Applications @ TidalScale
- Oracle ACE Director Alum
- Oracle Educator
  - Curriculum author and primary instructor, Oracle Program, University of Washington 1998-2009
  - Consultant: Harvard University
    - Guest lecturer at universities in Canada, Chile, Costa Rica, New Zealand, Norway, Panama
    - Frequent lecturer at Oracle conferences … 43 countries since 2008
- IT Professional
  - 2019 will be my 50th year in IT
  - First computer: IBM 360/40 in 1969: Fortran IV
  - Oracle Database since 1988-9 and Oracle Beta tester
  - The Morgan behind www.morganslibrary.org
  - Member Oracle Data Integration Solutions Partner Advisory Council
  - Founding member International TidalScale User Community (ITUC)

TidalScale™

# My Personal Website

www.morganslibrary.org

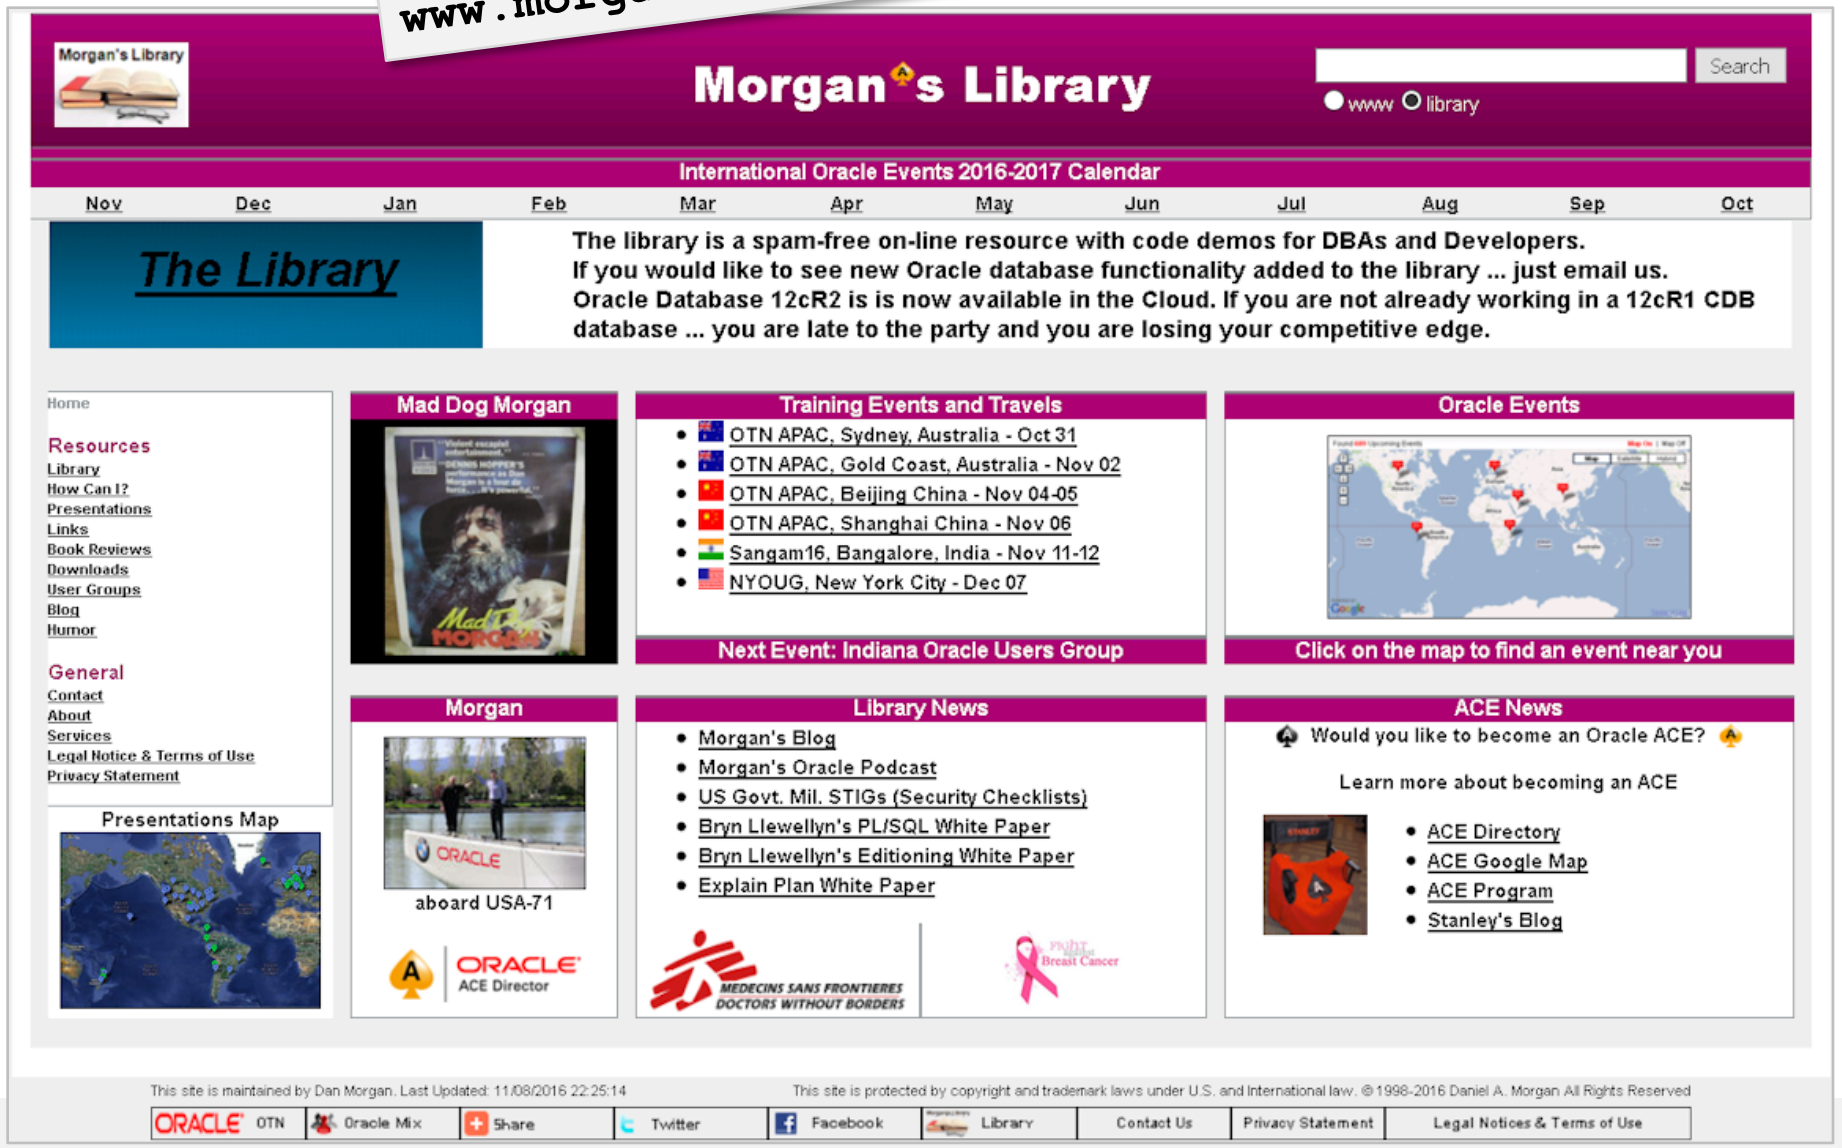


TidalScale™

## Performance Problems Have Serious Consequences . . .

- Internal and External customers have expectations
- There is a long history of disappoint
- Thus, we have Service Level Agreements

## When We Fail To Deliver . . .

- Internal customers develop their own solutions
- External customers go elsewhere
- SLA violations result in financial penalties
- Management wonders whether we are providing value

# Only 2 Things Matter In Business Computing . . .

## QoS

- Stability
- Security
- Scalability
- Usability
- Performance

## TCO

- Affordability

# The History of Oracle Performance Tuning . . .



TidalScale™

# How Many Books Read?

# How Many Oracle Tools Deployed?

- DBMS_SUPPORT (version 7.2)
- DBMS_TRACE (version 8.1.5)
- DBMS_MONITOR (version 10gR1)
- Oracle Enterprise Manager (OEM)
- StatsPack, ADDM, ASH, AWR, TKPROF, ....

```
ALTER SESSION SET tracefile_identifier = 'test_plan1';

ALTER SESSION SET EVENTS '10053 trace name context forever, level 1';

ALTER SESSION SET EVENTS '10046 trace name context forever, level 12';

-- execute SQL

ALTER SESSION SET EVENTS '10053 trace name context OFF';
ALTER SESSION SET EVENTS '10046 trace name context OFF';
or
ALTER SESSION SET SQL_TRACE=FALSE;

review the trace file in $ORACLE_BASE/diag/orabase/orabase/trace
```


TidalScale™

# How Many Tools Have You Purchased?




TidalScale™

# How Many Tools Have You Purchased?

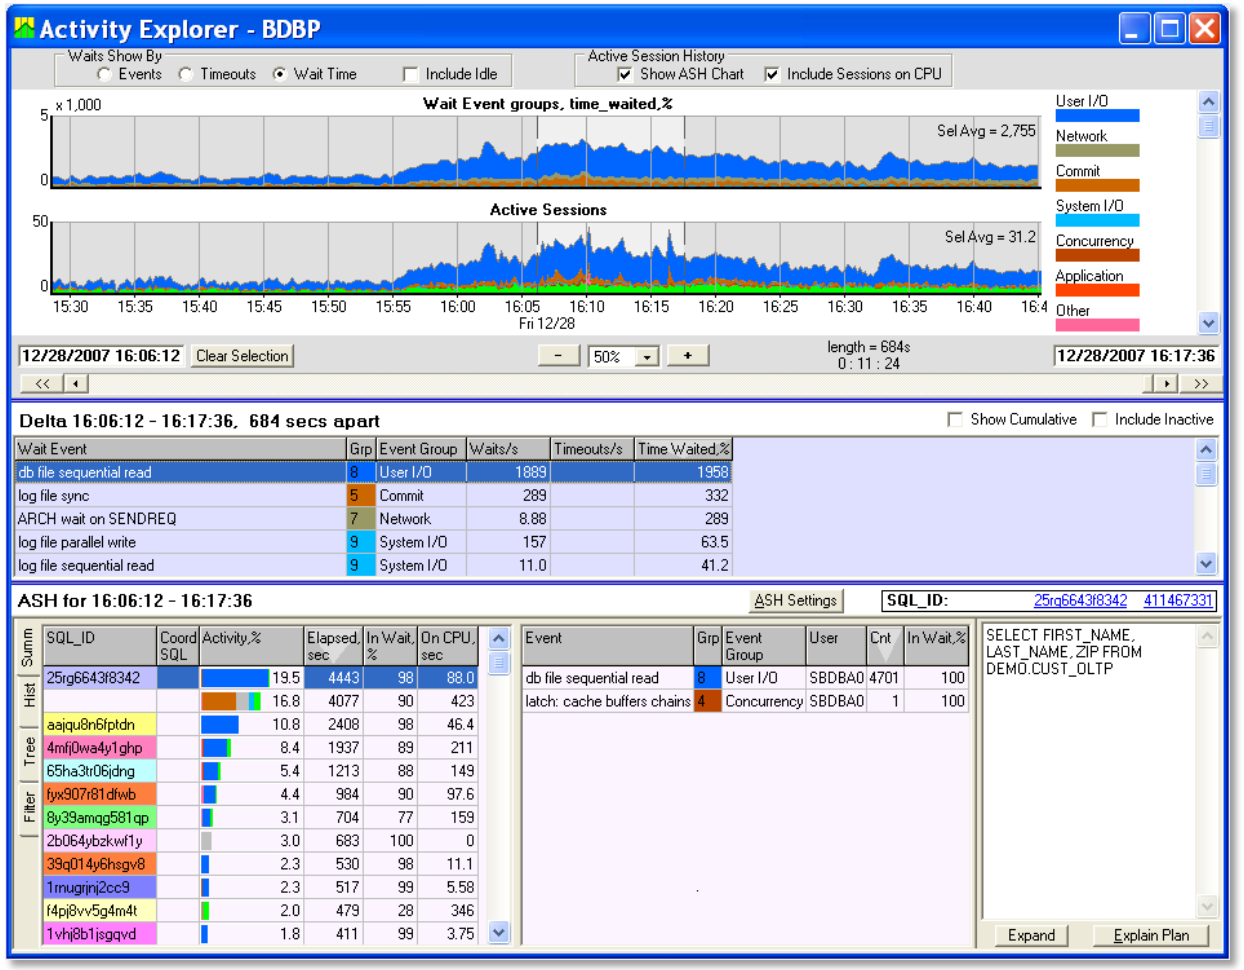




TidalScale™

# How Many Tools Have You Purchased?

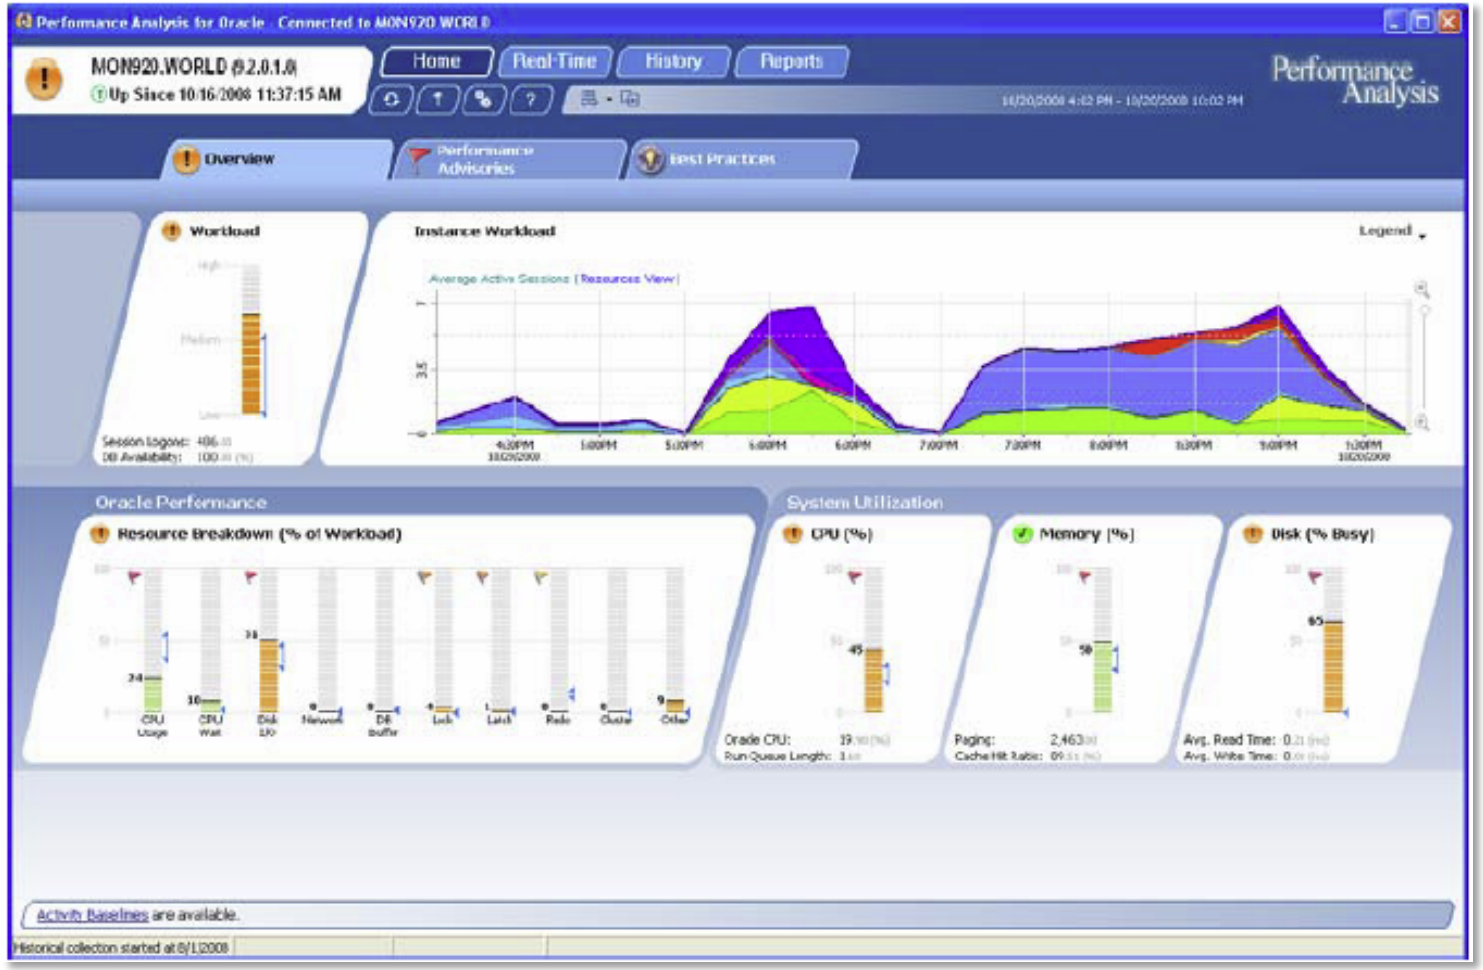




TidalScale™

# How Many Tools Have You Purchased?

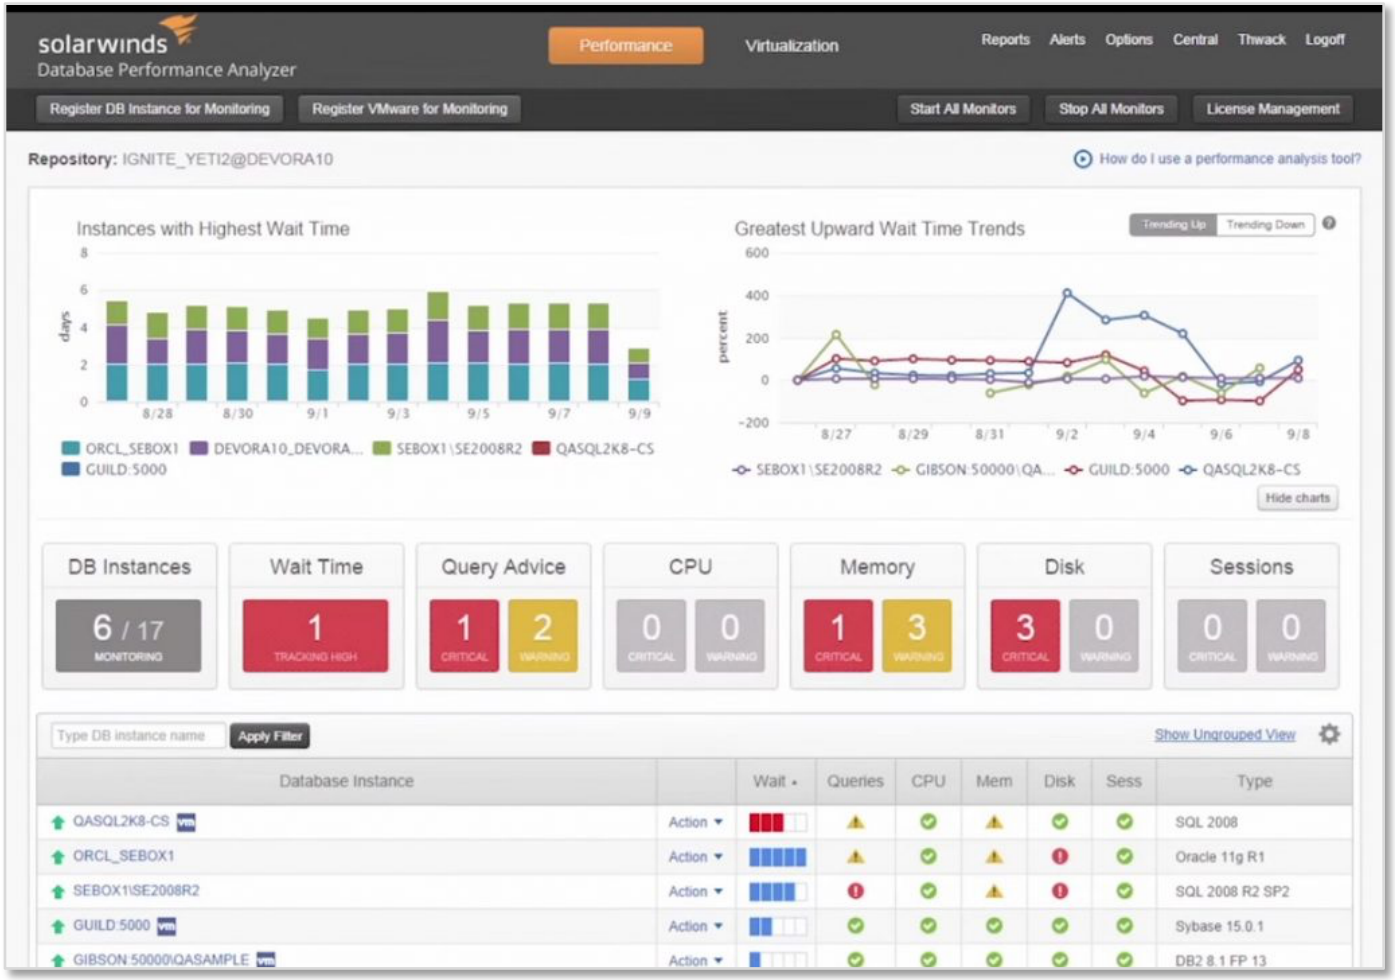




TidalScale™

# How Many Tools Have You Purchased?

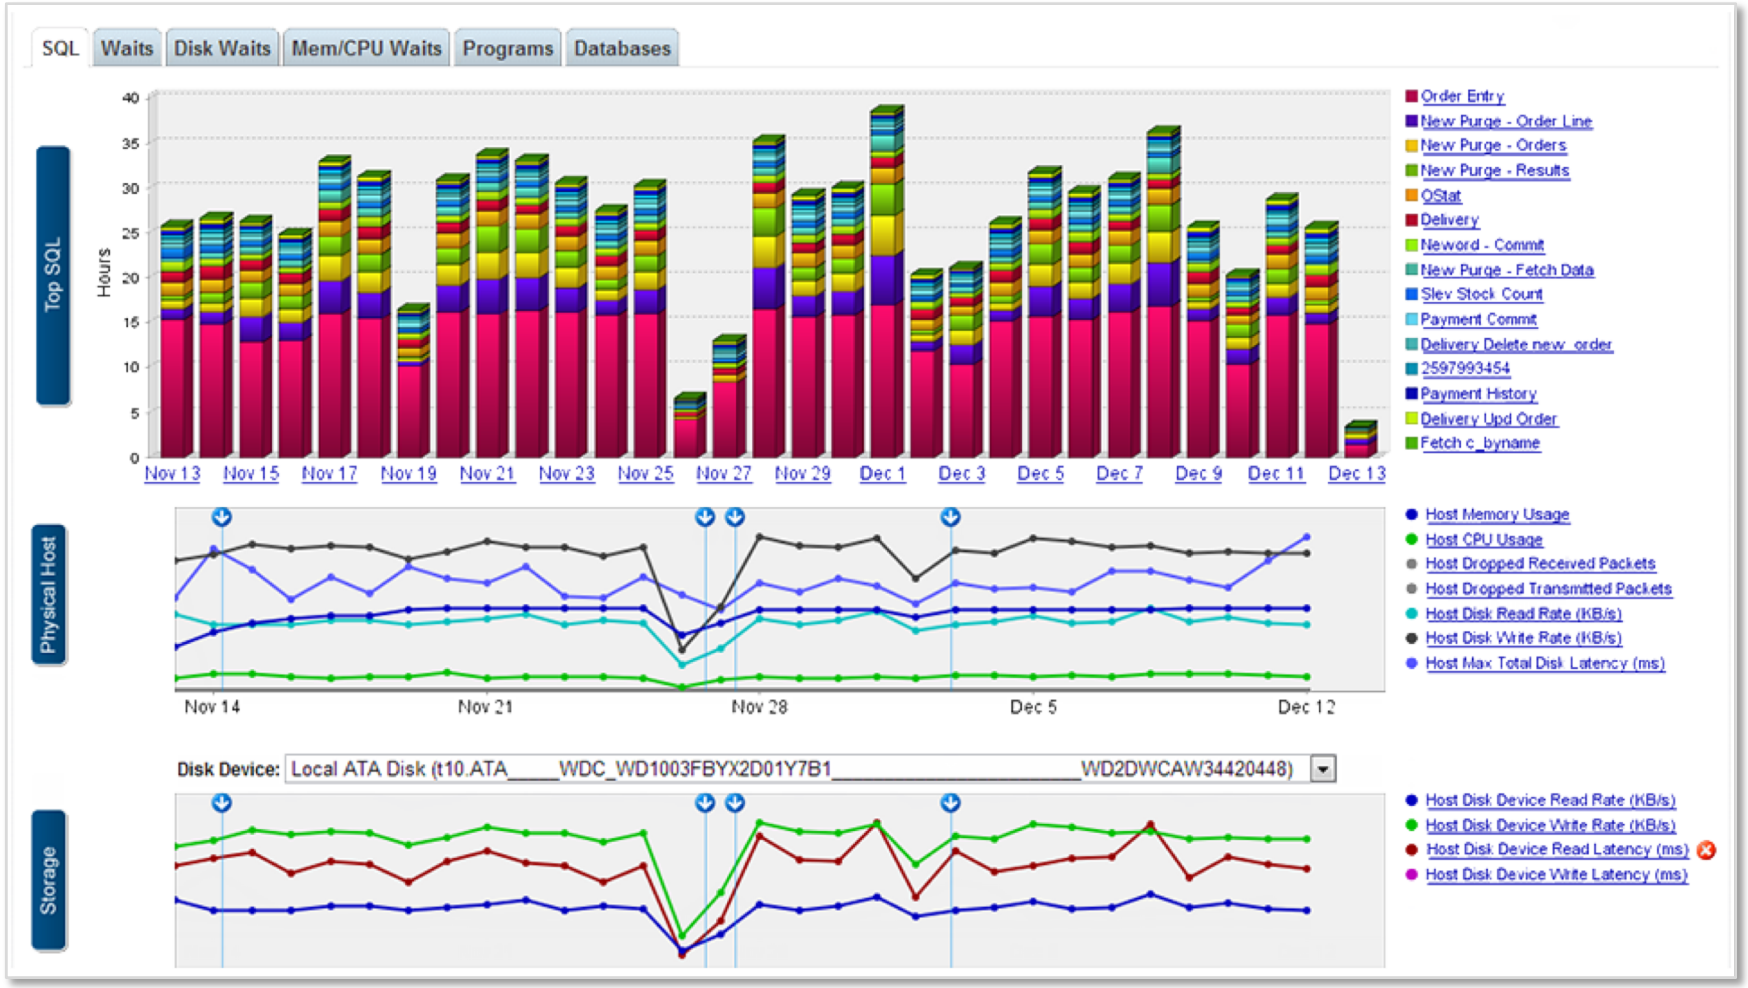


TidalScale™

# How Many Startup Parameters Configured?

| Initialization Parameter | Description |
|---|---|
| BITMAP_MERGE_AREA_SIZE | Specifies the amount of memory used to merge bitmaps retrieved from an index range scan |
| DB_BIG_TABLE_CACHE_PERCENT_TARGET | Specifies the cache section target size for automatic big table caching, as a percentage of the buffer cache |
| DB_nK_CACHE_SIZE | Holds 8K table and index blocks |
| CREATE_BITMAP_AREA_SIZE | Memory allocated for bitmap creation a larger value may speed up index creation |
| DB_BLOCK_BUFFERS | Specifies the number of database buffers in the buffer cache |
| DB_CACHE_SIZE | Specifies the size of the DEFAULT buffer pool for buffers with the primary block size |
| DB_FLASH_CACHE_SIZE | Specifies the size of the Database Smart Flash Cache |
| DB_KEEP_CACHE_SIZE | Specifies the size of the KEEP buffer pool |
| DB_RECYCLE_CACHE_SIZE | Specifies the size of the RECYCLE buffer pool |
| HASH_AREA_SIZE | Specifies the maximum amount of memory, in bytes, to be used for hash joins |
| JAVA_MAX_SESSIONSPACE_SIZE | Memory that holds Java state from one database call to another |
| JAVA_POOL_SIZE | Pool, from which the Java memory manager allocates most Java state during runtime execution |
| LARGE_POOL_SIZE | Specifies (in bytes) the size of the large pool allocation heap |
| LOG_BUFFER | Memory used when buffering redo entries to a redo log file |
| MEMOPTIMIZE_POOL_SIZE | Specifies the size of the memoptimize pool, a memory area in the SGA used by the Memoptimized Rowstore |

TidalScale™

| Initialization Parameter | Description |
|---|---|
| MEMORY_MAX_TARGET | Specifies the maximum value to which a DBA can set the MEMORY_TARGET initialization parameter |
| MEMORY_TARGET | Specifies the Oracle system-wide usable memory |
| OBJECT_CACHE_MAX_SIZE_PERCENT | specifies the percentage of the optimal cache size that the session object cache can grow past the optimal size |
| OBJECT_CACHE_OPTIMAL_SIZE | Specifies the size by which the session object cache is reduced when the cache size exceeds the maximum size |
| OLAP_PAGE_POOL_SIZE | Specifies the size of the OLAP page pool |
| PGA_AGGREGATE_LIMIT | Specifies a limit on the aggregate PGA memory consumed by the instance |
| PGA_AGGREGATE_TARGET | Specifies the target aggregate PGA memory available to all server processes attached to the instance |
| PRE_PAGE_SGA | Specifies whether Oracle reads the entire SGA into memory at startup so that O/S page table entries are pre-built for the SGA |
| SGA_MAX_SIZE | Specifies the maximum size of the SGA for the lifetime of the instance |
| SGA_MIN_SIZE | Specifies the minimum size of the SGA for the lifetime of the instance |
| SGA_TARGET | Specifies the total size of all SGA components |
| SHARED_POOL_RESERVED_SIZE | Specifies the shared pool space reserved for large contiguous requests for shared pool memory |
| SHARED_POOL_SIZE | Specifies the size of the shared pool which contains shared cursors, stored procedures, control and other structures |
| SORT_AREA_RETAINED_SIZE | Specifies the maximum amount of the user global area (UGA) memory retained after a sort run completes |
| SORT_AREA_SIZE | Specifies the maximum amount of memory Oracle will use for a sort |
| STREAMS_POOL_SIZE | Specifies the memory allocated for Streams, GoldenGate Integrated Capture and other related processes |
| USE_LARGE_PAGES | Specify the management of the database's use of large pages for SGA memory |

TidalScale™

AWR_report.html

# WORKLOAD REPOSITORY report for

| DB Name | DB Id | Unique Name | Role | Edition | Release | RAC | CDB |
|---|---|---|---|---|---|---|---|
| ORCL | 1499046141 | orcl | PRIMARY | EE | 12.2.0.1.0 | NO | NO |

| Instance | Inst Num | Startup Time |
|---|---|---|
| oracle | 1 | 25-Aug-18 16:08 |

| Host Name | Platform | CPUs | Cores | Sockets | Memory (GB) |
|---|---|---|---|---|---|
| oracle7002 | Linux x86 64-bit | 36 | 36 | 36 | 1153.16 |

| | Snap Id | Snap Time | Sessions | Cursors/Session |
|---|---|---|---|---|
| Begin Snap: | 2713 | 27-Aug-18 00:46:47 | 47 | .8 |
| End Snap: | 2714 | 27-Aug-18 00:58:57 | 103 | .8 |
| Elapsed: | | 12.18 (mins) | | |
| DB Time: | | 138.66 (mins) | | |

## Report Summary

### Top ADDM Findings by Average Active Sessions

| Finding Name | Avg active sessions of the task | Percent active sessions of finding | Task Name | Begin Snap Time | End Snap Time |
|---|---|---|---|---|---|
| Top SQL Statements | 11.40 | 70.40 | ADDM:1499046141_1_2714 | 27-Aug-18 00:46 | 27-Aug-18 00:58 |
| Undersized PGA | 11.40 | 3.47 | ADDM:1499046141_1_2714 | 27-Aug-18 00:46 | 27-Aug-18 00:58 |
| Undersized SGA | 11.40 | 2.82 | ADDM:1499046141_1_2714 | 27-Aug-18 00:46 | 27-Aug-18 00:58 |
| Unusual "Other" Wait Event | 11.40 | 2.27 | ADDM:1499046141_1_2714 | 27-Aug-18 00:46 | 27-Aug-18 00:58 |

### Memory Statistics

| | Begin | End |
|---|---|---|
| Host Mem (MB): | 1,180,832.7 | 1,180,832.7 |
| SGA use (MB): | 972,800.0 | 972,800.0 |
| PGA use (MB): | 361.9 | 7,848.7 |
| % Host Mem used for SGA+PGA: | 82.41 | 83.05 |

### Cache Sizes

| | Begin | End | | |
|---|---|---|---|---|
| Buffer Cache: | 96,768M | 96,768M | Std Block Size: | 8K |
| Shared Pool Size: | 202,163M | 202,149M | Log Buffer: | 495,048K |
| In-Memory Area: | 665,600M | 665,600M | | |

### Shared Pool Statistics

| | Begin | End |
|---|---|---|
| Memory Usage %: | 3.57 | 3.66 |
| % SQL with executions>1: | 91.10 | 90.41 |
| % Memory for SQL w/exec>1: | 89.90 | 87.87 |

## Foreground Wait Events

- s - second, ms - millisecond, us - microsecond, ns - nanosecon d
- Only events with Total Wait Time (s) >= .001 are shown
- ordered by wait time desc, waits desc (idle events last)
- %Timeouts: value of 0 indicates value was < .5%. Value of n ull is truly 0

| Event | Waits | %Time -outs | Total Wait Time (s) | Avg wait | Waits /txn | % DB time |
|---|---|---|---|---|---|---|
| direct path write temp | 25,156 | | 286 | 11.37ms | 613.56 | 3.44 |
| PGA memory operation | 191,325 | | 189 | .99ms | 4,666.46 | 2.27 |
| library cache: mutex X | 2,415 | | 27 | 11.24ms | 58.90 | 0.33 |

TidalScale™

# Full Database Caching . . .

- Primary use case: OLTP
  - Table scans and LOBs
- Disables the LRU algorithm so blocks in the Buffer Cache do not age out
- COMPATIBLE must be 12.0.0 or higher
- If using AMM (MEMORY_TARGET) or ASMM (SGA_TARGET) it is possible the buffer cache will resize, making the cache too small to hold the entire database
- Either size the memory parameters appropriately, or better still set the minimum size of the buffer cache by setting DB_CACHE_SIZE to an large value
- Objects are cached as accessed ... no pre-emptive loading
- LOBs defined as NOCACHE are cached when force full database cache is enabled
- Applies to the CDB and all PDBs when using multitenant
- To recover a control file force full database cache mode must be enabled

TidalScale™

# Enabling Full Database Caching . . .

```
conn / as sysdba

SHUTDOWN IMMEDIATE;
STARTUP MOUNT;
ALTER DATABASE force full database caching;
ALTER DATABASE OPEN;

SELECT force_full_db_caching FROM v$database;

FOR
---
YES
```

# In-Memory Column Store . . .

- Primary Use Case: DSS and DW where data will be accessed regularly for aggregation or analysis
- Data stored in row format in the block cache is also stored in a highly optimized columnar format that makes possible extreme compression
- In-Memory columnary storage is substantially more efficient than the HCC available in Exadata and ZFS
- Requires purchase of the Database In-Memory Option
- Data can be populated on-the-fly but is usually populated by identifying specific tables and columns

In-memory is probably the biggest performance booster out there, because its *not* just about putting things into memory - its about compression, restructuring the rows, using special CPU optimizations etc.
~ Connor McDonald ( Oracle ACE Director Alum, Oak Table Network)

- How are you going to put data in memory ... if you don't have enough DRAM?

TidalScale™

# Typically . . .

- Examples of Database In-Memory look like this ... a 4G store
- With TidalScale we build them big because we can

```
SQL> ALTER SYSTEM SET inmemory_size = 4G SCOPE=SPFILE;

SQL> shutdown immediate

Database closed.
Database dismounted.
ORACLE instance shut down.
SQL> startup
ORACLE instance started.

Total System Global Area 7516192768 bytes
Fixed Size                  3728304 bytes
Variable Size             838863952 bytes
Database Buffers         2365587456 bytes
Redo Buffers               13045760 bytes
In-Memory Area         4294967296 bytes
Database mounted.
Database opened.
```

- With TidalScale Software Defined Servers we can aggregate memory to 64TB
- We can create SGAs up to ~50TB
- An In-Memory Column Store can easily exceed the amount of available memory due to the high quality of columnar compression

www.dbta.com/Editorial/Trends-and-Applications/Tips-for-Accelerating-Analytics-with-Oracle-Database-In-Memory-110880.aspx

TidalScale™

## In-Memory Aggregation . . .

- Primary Use Case: DSS and DW
- **In-Memory Aggregation.** Designed especially for star schema query performance improvements
- Leverages a completely new set of optimizer methods collectively known as *vector transformation* to quickly identify the matching dimension keys for the query and then apply those limiting filters to the fact table
- Vector transformation leverages *single-instruction processing multiple data values* (SIMD) vector processing so that it's possible to read all of the possible values in one or more dimensional row sets *at one time*
- This reduces the need for complex bitmap join index structures to weed out non-matching dimension data
- SIMD vector processing can also benefit analytic functions that require aggregation – SUM, COUNT, AVERAGE, and many others

TidalScale™

## In-Memory Joins . . .

- Primary Use Case: DSS and DW
- **In-Memory Joins and Filtering**
- Joining data sources together – especially using hash joins when data sources are enormous – can consume equally huge amounts of memory and cpu
- IMCS leverages *Bloom Filters,* present "under the covers" in the Oracle optimizer since Oracle 10*g*R2, to convert the join to a filter that can be applied efficiently to a massive row sources reducing processing cycles
- IMCS has adapted and leveraged the Oracle Exadata concept of *storage indexes* to identify exactly where data isn't present and often completely avoids scanning vast sections of database objects when applying filters

TidalScale™

# This Is What 1TB Of Memory Looks Like To Most DBAs and CFOs

## WORKLOAD REPOSITORY COMPARE PERIOD REPORT

| Snapshot Set | DB Name | DB Id | Instance | Inst num | Release | Cluster | Host | Std Block Size |
|---|---|---|---|---|---|---|---|---|
| First (1st) | ORAP09 | 2886124853 | orap09 | 1 | 11.2.0.3.0 | NO | db20p03sh | 8192 |
| Second (2nd) | ORAP09 | 2886124853 | orap09 | 1 | 11.2.0.3.0 | NO | db20p03sh | 8192 |

| Snapshot Set | Begin Snap Id | Begin Snap Time | End Snap Id | End Snap Time | Avg Active Users | Elapsed Time (min) | DB time (min) |
|---|---|---|---|---|---|---|---|
| 1st | 33759 | 02-Apr-14 04:00:19 (Wed) | 33760 | 02-Apr-14 04:30:22 (Wed) | 0.8 | 30.1 | 24.5 |
| 2nd | 33807 | 03-Apr-14 04:00:13 (Thu) | 33808 | 03-Apr-14 04:30:16 (Thu) | 1.0 | 30.1 | 29.8 |
| *%Diff* | | | | | *22.2* | *0.0* | *21.6* |

### Host Configuration Comparison

| | 1st | 2nd | Diff | %Diff |
|---|---|---|---|---|
| Number of CPUs: | 80 | 80 | 0 | 0.0 |
| Number of CPU Cores: | 40 | 40 | 0 | 0.0 |
| Number of CPU Sockets: | 4 | 4 | 0 | 0.0 |
| Physical Memory: | 1031464.9M | 1031464.9M | 0M | 0.0 |
| Load at Start Snapshot: | 19.09 | 20.68 | 1.59 | 8.3 |
| Load at End Snapshot: | 11.49 | 11.18 | -.31 | -2.7 |
| %User Time: | 14.88 | 14.44 | -.44 | -3.0 |
| %System Time: | 1.08 | 1.05 | -.03 | -2.8 |
| %Idle Time: | 83.92 | 84.38 | .46 | 0.5 |
| %IO Wait Time: | .31 | .45 | .13 | 45.2 |

- 4 Sockets
- 40 cpu cores
- 80 threads
- 20 Oracle EE licenses ($950,000)
- 20 Diag & Tuning Licenses ($250,000)
- Total Licenses: $1,200,000
- Annual Support: $264,000

And all you get is 1TB of memory

TidalScale™

# TidalScale Metrics

- We used HammerDB to build three identical 500GB Oracle 12.2.0.1 databases on Oracle Enterprise Linux 7.4
  - Env 1: 512GB RAM … Oracle installation performed by OUI and DBCA <next><next><next>
  - Env 2: 1024GB RAM … A 2 node TidalScale pod with Database In-Memory enabled
- Nothing was customized
- The databases are 500GB and identical … created with the same script
- Adaptive Queries were not disabled … but we've tested both ways
- Evolving Baselines were not disabled … but we've tested both ways
- No Explain Plans were run
- No SQL Tuning was performed
- Every DML statement was generic TPC-H benchmark
- Are you ready to view the results?

# Test Results: Timing

Env 1: Run time – 13,249 seconds

Env 2: Run time – 558 seconds

A 25X performance improvement

TidalScale™

# Test Results: AWR Report (1:4)

**Env 1: Bare Metal**

| DB Name | DB Id | Unique Name | Role | Edition | Release | RAC | CDB |
|---|---|---|---|---|---|---|---|
| ORCL | 1512797244 | orcl | PRIMARY | EE | 12.2.0.1.0 | NO | NO |

| Instance | Inst Num | Startup Time |
|---|---|---|
| oracle | 1 | 24-Aug-18 15:08 |

| Host Name | Platform | CPUs | Cores | Sockets | Memory (GB) |
|---|---|---|---|---|---|
| oracle201 | Linux x86 64-bit | 32 | 16 | 2 | 503.81 |

| | Snap Id | Snap Time | Sessions | Cursors/Session |
|---|---|---|---|---|
| Begin Snap: | 42 | 26-Aug-18 05:29:27 | 82 | 1.0 |
| End Snap: | 43 | 26-Aug-18 05:41:24 | 78 | 1.0 |

**Env 2: TidalScale 2 Node Pod**

| DB Name | DB Id | Unique Name | Role | Edition | Release | RAC | CDB |
|---|---|---|---|---|---|---|---|
| ORCL | 1499046141 | orcl | PRIMARY | EE | 12.2.0.1.0 | NO | NO |

| Instance | Inst Num | Startup Time |
|---|---|---|
| oracle | 1 | 25-Aug-18 16:08 |

| Host Name | Platform | CPUs | Cores | Sockets | Memory (GB) |
|---|---|---|---|---|---|
| oracle7002 | Linux x86 64-bit | 36 | 36 | 36 | 1153.16 |

| | Snap Id | Snap Time | Sessions | Cursors/Session |
|---|---|---|---|---|
| Begin Snap: | 2713 | 27-Aug-18 00:46:47 | 47 | .8 |
| End Snap: | 2714 | 27-Aug-18 00:58:57 | 103 | .8 |
| Elapsed: | | 12.18 (mins) | | |
| DB Time: | | 138.66 (mins) | | |

TidalScale™

# Test Results: AWR Report (2:4)

**Env 1: Bare Metal**

**Load Profile**

| | Per Second | Per Transaction | Per Exec | Per Call |
|---|---|---|---|---|
| DB Time(s): | 12.4 | 442.6 | 2.14 | 0.32 |
| DB CPU(s): | 1.5 | 53.5 | 0.26 | 0.04 |
| Background CPU(s): | 0.0 | 0.3 | 0.00 | 0.00 |
| Redo size (bytes): | 7,014.7 | 251,425 .4 | | |
| Logical read (blocks): | 80,812.5 | 2,896,54 9.7 | | |
| Block changes: | 41.2 | 1,476.5 | | |
| Physical read (blocks): | 78,627.8 | 2,818,2 43.6 | | |
| Physical write (blocks): | 3,969.4 | 142,275.2 | | |
| Read IO requests: | 1,012.3 | 36,284.8 | | |
| Write IO requests: | 141.6 | 5,074.4 | | |
| Read IO (MB): | 614.3 | 22,017.5 | | |
| Write IO (MB): | 31.0 | 1,111.5 | | |
| IM scan rows: | 0.0 | 0.0 | | |
| Session Logical Read IM: | 0.0 | 0.0 | | |
| User calls: | 38.1 | 1,366.7 | | |
| Parses (SQL): | 4.9 | 174.2 | | |
| Hard parses (SQL): | 0.4 | 13.3 | | |
| SQL Work Area (MB): | 10.2 | 364.0 | | |
| Logons: | 0.2 | 6.1 | | |
| Executes (SQL): | 5.8 | 206.8 | | |
| Rollbacks: | 0.0 | 0.0 | | |
| Transactions: | 0.0 | | | |

**Env 2: TidalScale 2 Node Pod**

**Load Profile**

| | Per Second | Per Transaction | Per Exec | Per Call |
|---|---|---|---|---|
| DB Time(s): | 11.4 | 202.9 | 1.83 | 0.01 |
| DB CPU(s): | 11.2 | 199.8 | 1.80 | 0.01 |
| Background CPU(s): | 0.0 | 0.5 | 0.00 | 0.00 |
| Redo size (bytes): | 5,896.7 | 105,093.1 | | |
| Logical read (blocks): | 1,673,100.2 | 29,818,686.1 | | |
| Block changes: | 32.4 | 576.7 | | |
| Physical read (blocks): | 644.7 | 11,489.2 | | |
| Physical write (blocks): | 1,069.3 | 19,061.1 | | |
| Read IO requests: | 21.0 | 374.4 | | |
| Write IO requests: | 35.2 | 627.2 | | |
| Read IO (MB): | 5.0 | 89.8 | | |
| Write IO (MB): | 8.4 | 148.9 | | |
| IM scan rows: | 49,445,926.5 | 881,245,804.0 | | |
| Session Logical Read IM: | 1,669,664.9 | 29,757,459.9 | | |
| User calls: | 976.6 | 17,405.8 | | |
| Parses (SQL): | 4.5 | 80.2 | | |
| Hard parses (SQL): | 0.3 | 4.5 | | |
| SQL Work Area (MB): | 82.8 | 1,476.3 | | |
| Logons: | 1.0 | 18.1 | | |
| Executes (SQL): | 6.2 | 111.1 | | |
| Rollbacks: | 0.0 | 0.0 | | |
| Transactions: | 0.1 | | | |

- When you have sufficient memory logical reads replace far slower physical reads

TidalScale™

# Test Results: AWR Report (3:4)

**Env 1: Bare Metal**

**Memory Statistics**

| | Begin | End |
|---|---|---|
| Host Mem (MB): | 515,896.9 | 515,896.9 |
| SGA use (MB): | 155,136.0 | 155,136.0 |
| PGA use (MB): | 4,019.2 | 15,494.7 |
| % Host Mem used for SGA+PGA: | 30.85 | 33.07 |

**Env 2: TidalScale 2 Node Pod**

**Memory Statistics**

| | Begin | End |
|---|---|---|
| Host Mem (MB): | 1,180,832.7 | 1,180,832.7 |
| SGA use (MB): | 972,800.0 | 972,800.0 |
| PGA use (MB): | 361.9 | 7,848.7 |
| % Host Mem used for SGA+PGA: | 82.41 | 83.05 |

TidalScale™

# Test Results: AWR Report (4:5)

**Env 1: Bare Metal**

**Cache Sizes**

| | Begin | End | | |
|---|---|---|---|---|
| Buffer Cache: | 131,584M | 131,584M | Std Block Size: | 8K |
| Shared Pool Size: | 15,251M | 15,249M | Log Buffer: | 495,120K |
| In-Memory Area: | 0M | 0M | | |

**Env 2: TidalScale 2 Node Pod**

**Cache Sizes**

| | Begin | End | | |
|---|---|---|---|---|
| Buffer Cache: | 96,768M | 96,768M | Std Block Size: | 8K |
| Shared Pool Size: | 202,163M | 202,149M | Log Buffer: | 495,048K |
| In-Memory Area: | 665,600M | 665,600M | | |

- With TidalScale there is sufficient memory to put 500GB of segments into memory

TidalScale™

# Test Results: AWR Report (5:5)

## Segments by Physical Reads

**Env 1: Bare Metal**

- Total Physical Reads: 56,364,871
- Captured Segments account for 99.0% of Total
- When ** MISSING ** occurs, some of the object attributes may not be available

| Owner | Tablespace Name | Object Na me | Subobject Name | Obj. Type | Obj# | Dataobj# | Physical Reads | %Total |
|---|---|---|---|---|---|---|---|---|
| TPCH | TPCH500 | LINEITEM | | TABLE | 73611 | 73611 | 39,610,254 | 70.27 |
| TPCH | TPCH500 | ORDERS | | TABLE | 73604 | 73604 | 11,567,078 | 20.52 |
| TPCH | TPCH500 | PART | | TABLE | 73607 | 73607 | 4,361,323 | 7.74 |
| TPCH | TPCH500 | PARTSUPP | | TABLE | 73605 | 73605 | 272,313 | 0.48 |
| SYS | SYSAUX | WRH$_ACTIVE_SESSION_HISTORY | WRH$_ACTIVE_SESSION_HISTORY_1512797244_24 | TABLE PARTITION | 74028 | 74028 | 32 | 0.00 |

## Segments by Physical Reads

**Env 2: TidalScale 2 Node Pod**

- Total Physical Reads: 471,058
- Captured Segments account for 0.0% of Total
- When ** MISSING ** occurs, some of the object attributes may not be available

| Owner | Tablespace Name | Object Na me | Subobject Name | Obj. Type | Obj# | Dataobj# | Physical Reads | %Total |
|---|---|---|---|---|---|---|---|---|
| SYS | SYSAUX | WRH$_ACTIVE_SESSION_HISTORY | WRH$_ACTIVE_SESSION_HISTORY_1499046141_2706 | TABLE PARTITION | 98557 | 98557 | 113 | 0.02 |
| SYS | SYSAUX | SYS_LOB0000010641C00038$$ | | LOB | 10642 | 10642 | 32 | 0.01 |
| SYS | SYSAUX | WRH$_SQL_PLAN_PK | | INDEX | 10644 | 10644 | 11 | 0.00 |
| SYS | SYSAUX | WRH$_ENQUEUE_STAT_PK | | INDEX | 10676 | 10676 | 2 | 0.00 |
| SYS | SYSAUX | WRH$_SYSMETRIC_SUMMARY | | TABLE | 10843 | 10843 | 2 | 0.00 |

- TPCH tables are in memory so physical reads from TPCH500 are negligible

TidalScale™

# DBMS_INMEMORY_ADVISOR

- Released by Oracle Support and announced by Maria Colgan February, 2015
  - Compatible 11.2.0.3 and above
- Use to identify analytic workloads that will benefit from Database In-Memory
- Works by analyzing ASH and AWR data (diagnostic + tuning pack required)
- Download MOS note 1965343.1
- Oracle's tool will advise you on the value of DBIM
- TidalScale will advise you on the performance and security benefits of moving to Software Defined Servers

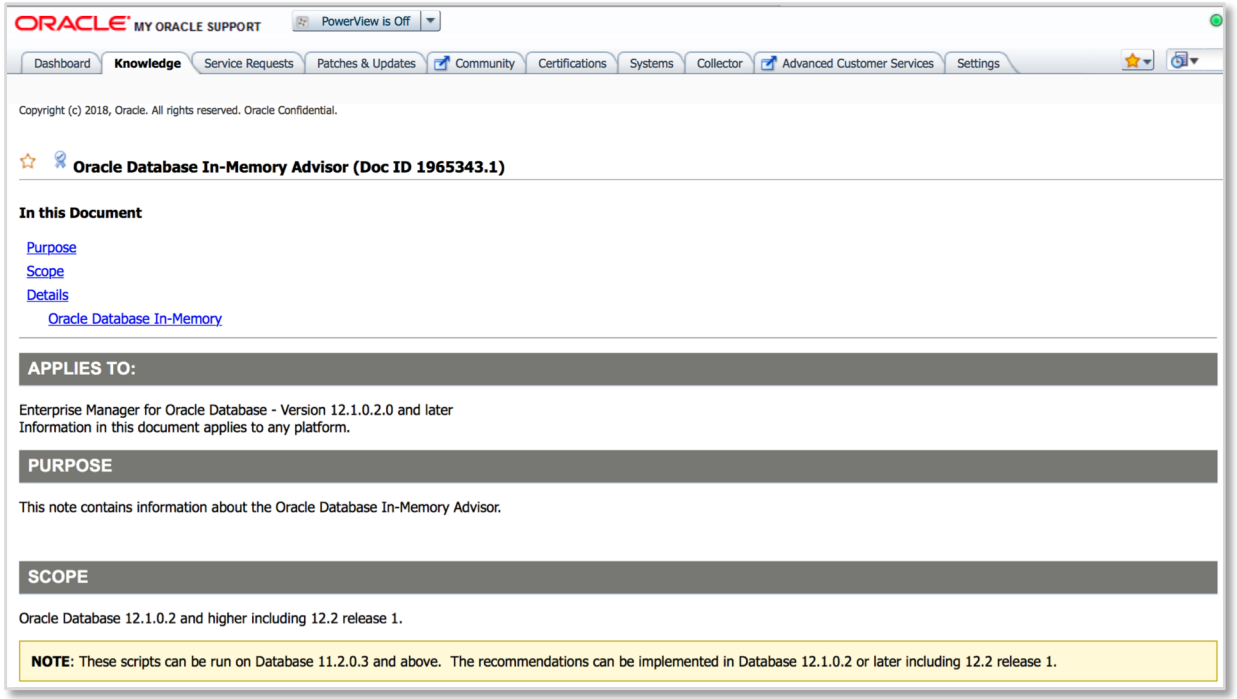


TidalScale™

## Wrap Up . . .

- Memory is 300X faster than flash
- The more of your data you can cache in memory the more your performance will improve due to reducing PIO and increasing LIO
- If you are running OLTP loads use FORCE FULL to disable the LRU algorithm
- If you are running DW or DSS loads they will dramatically benefit from Database In-Memory
  - To  not be limited to a few tables and a few columns
  - To avoid having to invest the time required to figure out which ones
  - Use TidalScale Software Defined Servers and build the right environment
- TidalScale Software Defined Servers let you create an ideal environment with the cpu and memory customized to
  - Improve performance
  - Control your licensing cost

TidalScale™

# Next Steps

**Contact me directly to**

- Answer questions about TidalScale Software Defined Servers
- Present TidalScale Software Defined Servers to your team
- Identify opportunities in your organization for Software Defined Servers

**daniel.morgan@tidalscale.com**

TidalScale™