



# Performance Diagnostics With AWR Reports

Richard Foote Consulting



# Richard Foote



- Working in IT for over 30+ years , 20+ years with Oracle Database
- 19 years employed in Australian Federal Government in various IT roles
- Responsible for many large scale, mission critical, “life-dependant” classified Oracle systems, tuned numerous databases often with 10x performance improvements
- Worked for Oracle Corporation between 1996 and 2002 and between 2011 and 2017
- In September 2017, started my own independent company Richard Foote Consulting
- Oracle OakTable Member since 2002 and awarded Oracle ACE Director in 2008
- Regular speaker at user group meetings and conferences such as Oracle OpenWorld, IOUG Collaborate, Hotsos Symposium, AUSOUG InSync, ODTUG Kscope, UKOUG Tech Conference, E4 Enkitec Extreme Exadata Expo, ...
- Richard Foote's Oracle Blog: <https://richardfoote.wordpress.com>
- Richard Foote Consulting: <https://richardfooteconsulting.com>
- Spend as much free time as possible listening to the music of David Bowie !!





Richard Foote's Oracle Blog  
Focusing Specifically On Oracle Indexes, Database Administration and Some Great Music

home richard foote presentations & demos index internals seminar public appearances recommendations

## Introduction To Reverse Key Indexes: Part III (A Space Oddity)

January 18, 2008

Posted by Richard Foote in [Oracle Indexes](#).  
[9 comments](#)

A possibly significant difference between a Reverse and a Non-Reverse index is the manner in which space is used in each index and the type of block splitting that takes place.

Most Reverse Key Indexes are created to resolve contention issues as a result of monotonically increasing values. As monotonically increasing values get inserted, each value is greater than all previous values (providing there are no outlier values present) and so fill the "right-most" leaf block. If the "right-most" block is filled by the maximum current value in the index, Oracle performs 90-10 block splits meaning that full index blocks are left behind in the index structure. Assuming no deletes or updates, the index should have virtually 100% used space.

However, it's equivalent Reverse Key index will have the values reversed and dispersed evenly throughout the index structure. As index blocks fill, there will be a very remote chance of it being due to the maximum indexed value and 50-50 block splits will result. The PCT\_USED is likely therefore to be significantly less, averaging approximately 70-75% over time.

Therefore, for indexes with no deletions, a Reverse Key index is likely to be less efficient from a space usage point of view.

However, if there are deletions, the story may differ.

Deleted space can be reused if an insert is subsequently made into an index block with deleted entries or if a leaf block is totally emptied. However, if a leaf block contains any non-deleted entries and if subsequent inserts don't hit the leaf block, then the deleted space can not be reused. As monotonically increasing values in a non-reverse index only ever insert into the "right-most" leaf block, it won't be able to reuse deleted space if leaf blocks are not totally emptied. Overtime, the number of such "almost but not quite empty" index leaf blocks may in some scenarios increase to significant levels and the index may continue to grow at a greater proportional rate than the table (where the reuse of space is set and controlled by the PCTUSED physical property).

search  go!

### Contact Details

If you wish to contact me directly, please do so at [richard.foote@bigpond.com](mailto:richard.foote@bigpond.com)

### Recent Posts

- Introduction To Reverse Key Indexes: Part III (A Space Oddity)
- European Summer Tour and Other Public Appearances
- Introduction To Reverse Key Indexes: Part II (Another Myth Bites The Dust)
- Introduction To Reverse Key Indexes: Part I
- Introduction to Fake / Virtual / NOSEGMENT indexes
- 8 Things You May Not Know About Indexes
- 10,000 Hits Already !!
- Introduction To Linguistic Indexes - Part II
- DBMS\_STATS METHOD\_OPT default behaviour changed in 10g. Be careful ...
- Introduction To Linguistic Indexes - Part I
- Differences between Unique and Non-Unique Indexes (Part III)
- Merry Christmas and a Happy Index Rebuild Free New Year !!
- Differences between Unique and Non-Unique Indexes (Part II)
- Local Index Issue With Partitioned PK and Unique Key Constraints
- Do ROWID Index Row Entry Columns Impact Index Block Splits ?

### Recent Comments

Richard Foote

# Oracle Indexing Internals and Best Practices 5 Day Webinars



**8-12 October, 6-10 November 2018 (4 Hours Daily)**

*Of benefit to DBAs, Developers, Solution Architects and anyone else interested in designing, developing or maintaining high performance Oracle-based applications/databases.*

Examines most available index structures/options & discusses in considerable detail how indexes function, how/when they should be used & how they should be maintained. Also how indexes are costed & evaluated by the Cost Based Optimizer (CBO) & how appropriate data management practices are vital for an effective indexing strategy. Covers many useful tips and strategies to maximise the benefits of indexes on application/database performance & scalability.

Running between **10am-2pm Zurich Time 8<sup>th</sup> – 12<sup>th</sup> October 2018**

**[richardfooteconsulting.com/indexing-webinar/](http://richardfooteconsulting.com/indexing-webinar/)**





# DB Running Slow, Where To Start ??

## WORKLOAD REPOSITORY report for

DB Name	DB ID	Instance	Inst num	Startup Time	Release	RAC
BOWIE	1234567890	BOWIE	1	25-Jun-99 12:56	12.1.0.2.0	NO

Host Name	Platform	CPUs	Cores	Sockets	Memory (GB)
bowie011	Linux x86 64-bit	6	8	8	31.36

	Snap ID	Snap Time	Sessions	Cursors/Session
Begin Snap:	4724	30-Jul-99 13:00:20	278	6.7
End Snap:	4725	30-Jul-99 14:00:21	333	7.0
Elapsed:		60.17 (mins)		
DB Time:		773.53 (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	12.85	79.67	ADDM:1234567890_1_4725	30-Jul-99 13:00:20	30-Jul-99 14:00:21
Row Lock Waits	12.85	56.91	ADDM:1234567890_1_4725	30-Jul-99 13:00:20	30-Jul-99 14:00:21
Unusual "Network" Wait Event	12.85	6.58	ADDM:1234567890_1_4725	30-Jul-99 13:00:20	30-Jul-99 14:00:21
Hard Parse Due to Liberal Usage	12.85	6.45	ADDM:1234567890_1_4725	30-Jul-99 13:00:20	30-Jul-99 14:00:21
Unusual "Network" Wait Event	12.85	3.75	ADDM:1234567890_1_4725	30-Jul-99 13:00:20	30-Jul-99 14:00:21

### Load Profile

	Per Second	Par Transaction	Par Exec	Par Call
DB Time(s):	12.9	0.7	0.00	0.00
DB CPU(s):	4.0	0.2	0.00	0.00
Background CPU(s):	0.1	0.0	0.00	0.00
Redo size (bytes):	216,162.4	12,514.4		
Logical read (blocks):	245,114.3	14,060.4		
Block changes:	1,102.0	63.2		
Physical read (blocks):	24,546.8	1,408.1		
Physical write (blocks):	62.5	3.0		
Read IO requests:	1,474.2	84.6		
Write IO requests:	23.0	1.3		
Read IO (MB):	191.6	11.0		
Write IO (MB):	0.4	0.0		
IM scan rows:	0.0	0.0		
Session Logical Read (M):				
User calls:	7,235.3	415.0		
Parses (SQL):	802.9	46.1		
Hard parses (SQL):	34.1	2.0		
SQL Work Area (MB):	30.7	1.8		
Logons:	0.6	0.0		
Executes (SQL):	4,933.3	283.0		
Rollbacks:	0.3	0.0		
Transactions:	17.4			

### Instance Efficiency Percentages (Target 100%)

Buffer Nowait %:	100.00	Redo Nowait %:	100.00
Buffer Hit %:	99.21	In-memory Sort %:	100.00
Library Hit %:	97.68	Sort Parse %:	95.75
Execute to Parse %:	83.73	Latch Hit %:	99.39
Parse CPU to Parse Elapsed %:	75.58	Non-Parse CPU:	63.72
Flash Cache Hit %:	0.00		

### Top 10 Foreground Events by Total Wait Time

Event	Waits	Total Wait Time (sec)	Wait Avg(ms)	% DB time	Wait Class
enq: TX - row lock contention	346	26.4K	76340.32	56.9	Application
DB CPU		14.9K		31.3	
SQL*Net message from dblink	17,686,859	3055.2	0.17	6.6	Network
SQL*Net more data from dblink	1,126,300	1738.4	1.54	3.7	Network
latch free	26,770	623.3	23.29	1.3	Other
direct path read	668,179	170.6	0.26	.4	User I/O
cursor: pin S wait on X	1,752	168.8	97.53	.4	Concurrency
log file sync	62,198	105.5	1.70	.2	Commit
library cache: mutex X	35,239	67.7	1.92	.1	Concurrency
db file sequential read	4,224,329	63.5	0.02	.1	User I/O

### Wait Classes by Total Wait Time



# Where's The Milk ?

- My wife kindly pops down to the local shop to buy some milk. A full 60 minutes later, she finally returns and plops the milk on the kitchen bench.
- This “**response time**” is clearly unacceptable, I needed the milk within 5 minutes...
- So how do I improve the response time ?

Note: In real-life, my wife would of course say go get the milk yourself !!





# Response Time

Response times can simplistically be broken up into two basic components:

Time it takes doing something

+

Time it takes waiting on something



**Response Time = Doing Time + Waiting Time**

To tune effectively and reduce response times, we need to focus on where most time is being spent !!



# Database Time (DB Time)

- Total time in database calls by **foreground sessions**
- Includes **CPU** time, **IO** time and **non-idle wait** time

***Database time is total time spent by user processes either actively working or actively waiting in a database call***

- ✓ Always encouraged when a DBA can state typical DB Time and Active Sessions of their databases



# Fundamental Concepts Summary



**Database Time (DB Time) =**

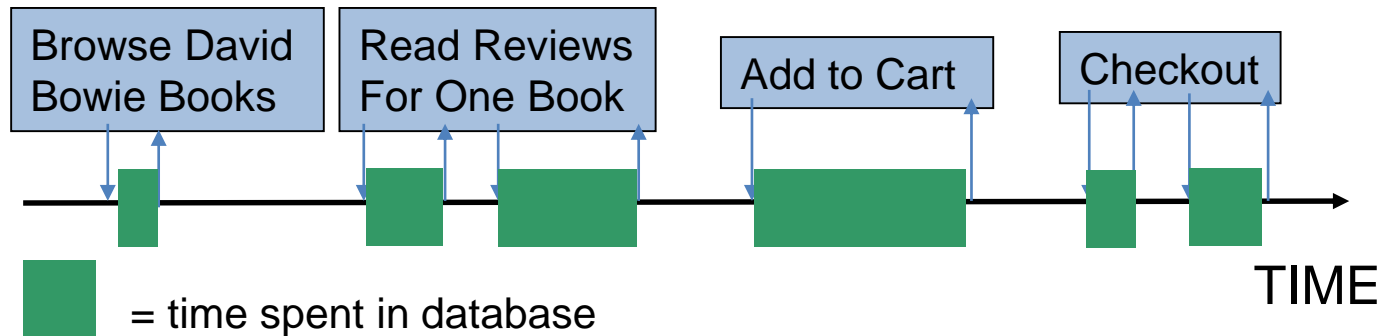
Total time foreground sessions spend in all database calls

**Active Session =**

Session currently spending time in a database call

**Average Activity of the Session (% Activity) =**

The ratio of time active to total wall clock time






# Active Sessions

DB Time = Sum of DB Time Over All Sessions

Avg. Active Sessions = Sum of Avg. Activity Over All Sessions

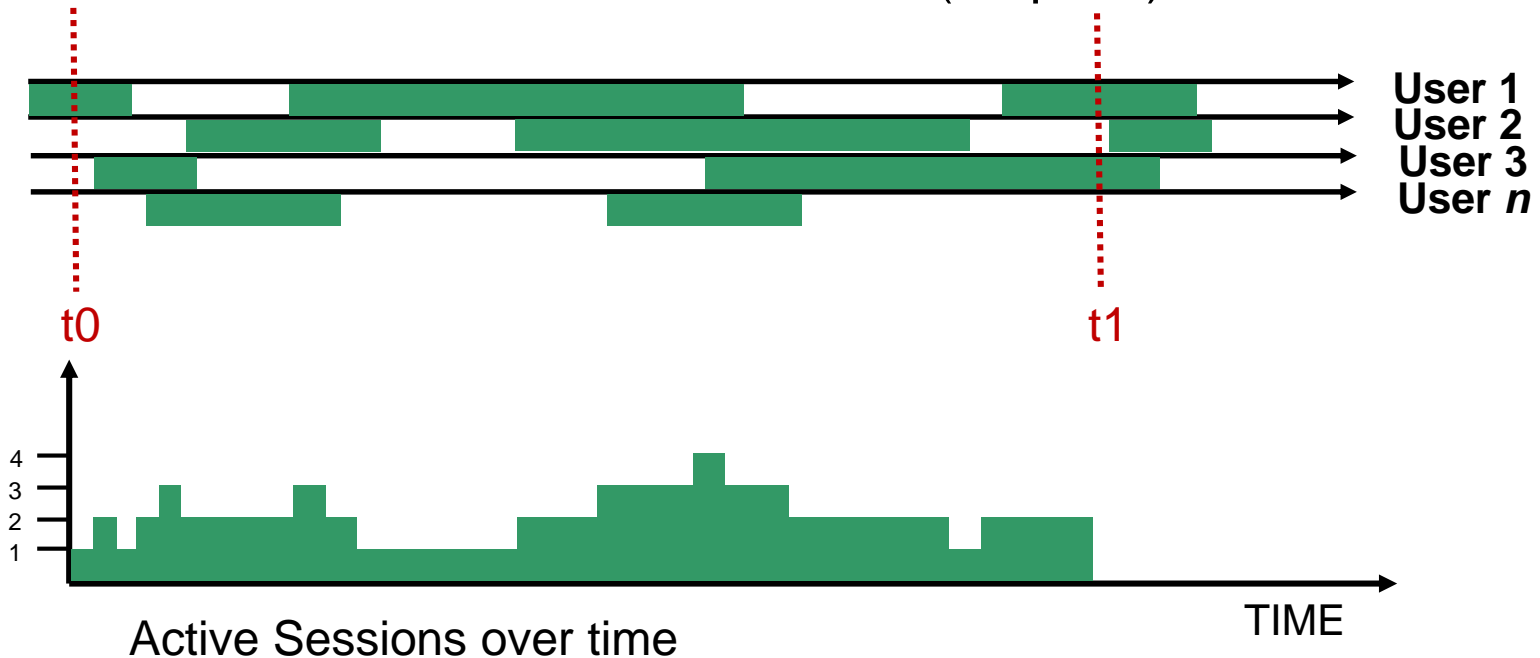
At time  $t$  we have 2 active sessions





# Visualizing DB Time

$$\text{Avg. Active Sessions} = \frac{\text{Total Database Time}}{\text{Wall Clock (Elapsed) Time}}$$



# Host Performance and DB Time



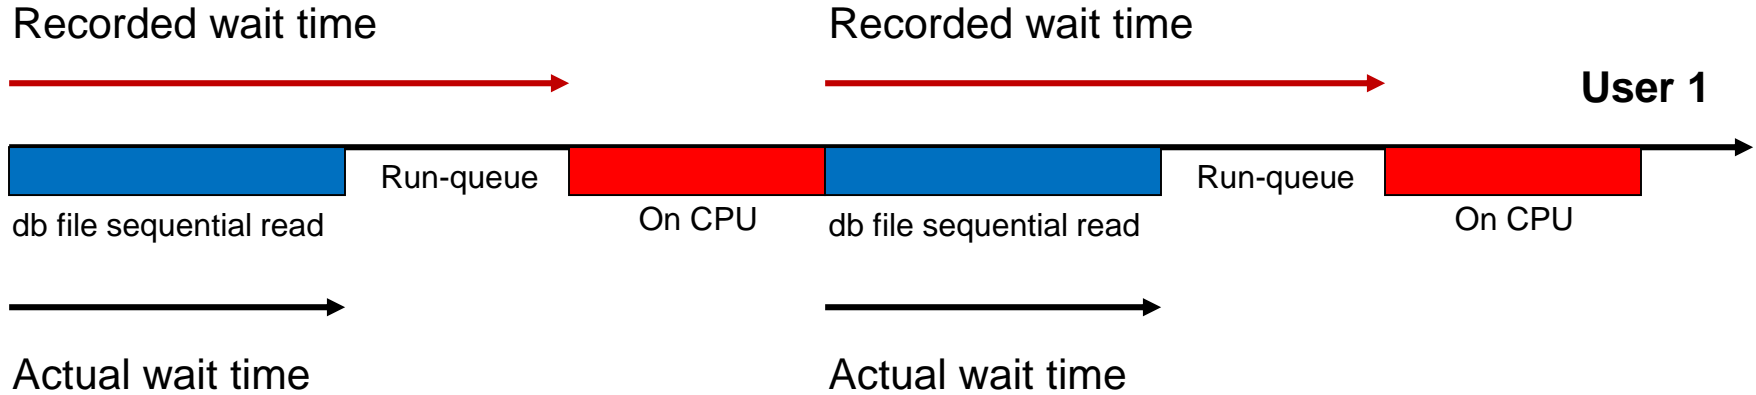
- Host is CPU-bound
  - => foregrounds accumulate active run-queue time
  - => wait event times are artificially inflated
  - => DB time increases



***Tune for CPU before waits when CPU constrained***



# CPU Run-queue and DB Time



***DB time can be inflated when host is CPU-bound***

# If Database Performance Is Poor Globally



- Everyone (or significant number) experiencing some form of performance degradation (response times have increased)
- Overall DB time can be expected to have increased
- Can potentially use session level analysis as specific sessions can be indicative of general issues
- But a system-wide view of the database can be extremely beneficial
- Enter **Automatic Workload Repository (AWR) report**
- Provides a detailed analysis of where/how all DB time is generated

# Automatic Workload Repository (AWR)



- Built-in automatic in-memory performance statistics
- MMON background process writes performance statistics snapshots to disk every hour (default)
- Snapshots kept for 8 days (default)
- AWR Report provides a database wide report based on these metrics for given snapshot period
- Automatic Database Diagnostics Monitor (ADDM) finds top “problems”
- Requires Diagnostics Pack (or High Performance Oracle DB Cloud Service)



# How To Generate AWR Report

Parameter: CONTROL\_MANAGEMENT\_PACK\_ACCESS=diagnostic

@\$ORACLE\_HOME/rdbms/admin/awrrpt.sql (for local database)

Enter value for report\_type: html (or text or active-html)

Enter value for num\_days: 1

Enter value for begin\_snap: 42

Enter value for end\_snap: 43 (**Note: Correct scoping is vital !!**)

Enter value for report\_name: bowie\_report



# How To Generate AWR Reports



@\$ORACLE\_HOME/rdbms/admin/awrrpti.sql (for specific database)

@\$ORACLE\_HOME/rdbms/admin/awrgrpt.sql (for local RAC database)

@\$ORACLE\_HOME/rdbms/admin/awrgrpti.sql (for specific RAC database)

@\$ORACLE\_HOME/rdbms/admin/awrsqrpt.sql (for specific SQL Id on local database)

@\$ORACLE\_HOME/rdbms/admin/awrsqrpi.sql (for specific SQL Id on specific database)



# Having A Baseline Can Help

- Having a baseline of when the database is hunky dory helps when things go bad:

```
BEGIN DBMS_WORKLOAD_REPOSITORY.create_baseline (  
    start_snap_id => 420,  
    end_snap_id => 422,  
    baseline_name => 'HUNKY DORY');  
END; /
```

- If the database is suddenly problematic, then DB Times have increased. Question: Why ?
- Easily seeing what has “changed” can help pinpoint introduced issues
- Can generate AWR Diff report that details differences between 2 AWR reports:
  - @\$ORACLE\_HOME/rdbms/admin/awrddrpt.sql (difference report for local database)
  - @\$ORACLE\_HOME/rdbms/admin/awrgdrpt.sql (difference report for RAC database)

# Useful AWR scripts



By default, AWR reports are generated every hour and retained for 8 days. To change:

```
BEGIN DBMS_WORKLOAD_REPOSITORY.modify_snapshot_settings(  
    retention => 86400, -- minutes, 60 Days  
    interval => 30, -- minutes  
    topnsql => 50);  
END;
```

To manually create a new AWR snapshot:

```
EXEC DBMS_WORKLOAD_REPOSITORY.create_snapshot;
```

To drop a range of unwanted snapshots:

```
BEGIN DBMS_WORKLOAD_REPOSITORY.drop_snapshot_range (  
    low_snap_id => 42, high_snap_id => 84);  
END;
```

# Oracle Enterprise Manager or SQL Developer



The screenshot displays the Oracle SQL Developer interface with a 'WORKLOAD REPOSITORY report for' window open. The report provides detailed performance metrics for the database instance 'BOWIE'.

DB Name	DB Id	Instance	Inst num	Startup Time	Release	RAC
BOWIE	1427541897	bowie	1	26-Oct-15 12:15	12.1.0.2.0	NO

Host Name	Platform	CPUs	Cores	Sockets	Memory (GB)
RFOOTE-AU	Microsoft Windows x86_64-bit	4	2	1	7.70

Snap Id	Snap Time	Sessions	Carblocks/session
Begin Snap	419 26-Oct-15 12:27:53	36	1.6
End Snap	422 26-Oct-15 15:00:27	39	1.5
Elapsed	182.88 (mins)		
DB Time	8.45 (mins)		

### Report Summary

#### Load Profile

	Per Second	Per-Transaction	Per Exec	Per Call
DB Time(s)	0.0	0.1	0.00	0.00
DB CPU(s)	0.0	0.0	0.00	0.00
Background CPU(s)	0.0	0.1	0.00	0.00
Redo size (bytes)	7,832.7	150,757.5		
Logical read (blocks)	222.5	5,476.3		
Block changes	50.8	1,250.0		
Physical read (blocks)	96.1	2,364.0		

# Focus Today On “How” To Best Interpret AWR Reports



The screenshot shows an Oracle AWR (Automatic Workload Repository) report for a database instance named BOWIE. The report is displayed in a web browser window with the following sections:

### WORKLOAD REPOSITORY report for

DB Name	DB ID	Instance	Inst name	Database	SID	Host
BOWIE	3057185373	BOWIE	1	10.2.0.5.0	10	jenky-d01

	Snap Id	Snap Time	Sessions	Cursors/Session
Begin Snap	32746	15-Feb-14 14:30:40	24	5.5
End Snap	32748	15-Feb-14 15:01:30	26	9.4
Elapsed		00:30 (mins)		
DB Time		47:02 (mins)		

### Report Summary

#### Cache Sizes

	Begin	End	
Buffer Cache	9,729M	9,729M	Cell Block Size
Shared Pool Size	432M	432M	Log Buffer

#### Load Profile

	Per Second	Per Transaction
Block gets	344,731.26	174,607.52
Logical reads	23,228.95	7,843.25
Block changes	2,324.20	762.50
Physical reads	94.41	33.66
Physical writes	52.04	13.51
User calls	3.51	8.71
Paras	61.06	16.75
Hard paras	3.09	3.02
Sorts	13.81	3.88
Logons	3.00	3.01
Discards	1,744.75	872.06
Transactions	3.00	

% Blocks changed per Read	0.63	Recursive Call %	99.57
Parallel per transaction %	0.00	Rows per Sort	29.09

#### Instance Efficiency Percentages (Target 100%)

Buffer Inwait %	100.00	Parse InWait %	100.00
Buffer Hit %	99.87	in-memory Sort %	100.00
Library Hit %	99.99	Call Parse %	99.99
Commit to Parse %	97.21	Latch Hit %	100.00
Parse CPU to Parse Elapsed %	95.42	% Non-Parse CPU	99.62

#### Shared Pool Statistics

	Begin	End
--	-------	-----



# Example: Log File Parallel Write vs. Log File Sync Times

Average log file parallel write times:

$100 \times 1\text{ms}, 20 \times 9\text{ms}$  and  $1 \times 809\text{ms} = 100+180+809/111 = 1089/121 = 9\text{ms}$

Average log file sync times (very very *simplistically*):

Assuming fewer log file syncs per log write over 1ms on average, waiting 1/2 the actual LFPW time:

$(200 \times 1 + 100 \times 4.5 + 30 \times 405)/(200 + 100 + 30) = 12800/330 = 39\text{ms}$

Note: In reality, many more transactions on average likely impacted during slower LFPW periods...

# What If Only Some Sessions Have Performance Issues?



- In many scenarios, “overall” database performance appears fine
- However, a subset of users are complaining about performance
- Looking at database wide metrics are useless as useful data is drowned out by general system info
- Once data is aggregated, specific information is lost
- Averages can be very misleading (database/session)
- Database level stats don’t necessarily translate to specific performance issue
- Need to be able to drill down on the DB time of just the specific users
- Enter **Active Session History (ASH)**
- Can also use **Extended SQL Trace** (DBMS\_MONITOR package)

# TKPROF: Standard SQL Trace



call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.00	0.00	0	6	0	0
Execute	1	0.06	62.46	0	7772	4	1
Fetch	0	0.00	0.00	0	0	0	0
total	2	0.06	62.47	0	7778	4	1

Rows (1st)	Rows (avg)	Rows (max)	Row Source Operation
0	0	0	UPDATE ZIGGY (cr=7772 pr=0 pw=0 time=62462379 us)
1	1	1	TABLE ACCESS FULL ZIGGY (cr=7771 pr=0 pw=0 time=63686 us cost=2122 size=20 card=1)

Classic example: why the massive difference between CPU time and elapsed time ?



# Let's Check Database Level AWR Report



## Top 5 Timed Events

```
~~~~~
```

Event	Waits	Time (s)	% Total Ela Time
db file sequential read	943,457	18,678	46.33
db file scattered read	381,532	6,059	15.03
CPU time		5,627	13.96
direct path read (lob)	326,048	5,550	13.77
SQL*Net more data to client	204,957	3,051	7.57

It certainly looks like we might have an I/O related problem here ...

# TKPROF: Extended SQL Trace



Elapsed times include waiting on following events:

Event waited on	Times	Max. Wait	Total waited
-----	waited	-----	-----
enq: TX - row lock contention	1	60.86	60.86
SQL*Net message to client	1	0.00	0.00
SQL*Net message from client	1	5.90	5.90

\*\*\*\*\*

With extended SQL wait data, we're lead to the actual cause of the problem

Note: Targeted ASH data is generally sufficient to also determine actual problem

## Introduction To Reverse Key Indexes: Part III (A Space Oddity) January 18, 2008

Posted by Richard Foote in [Oracle Indexes](#).  
[9 comments](#)

A possibly significant difference between a Reverse and a Non-Reverse index is the manner in which space is used in each index and the type of block splitting that takes place.

Most Reverse Key Indexes are created to resolve contention issues as a result of monotonically increasing values. As monotonically increasing values get inserted, each value is greater than the previous value (providing there are no outlier values present) and so the "right-most" leaf block. If the "right-most" block is filled by the maximum current value in the index, Oracle performs 90-10 block splits meaning the full index blocks are left behind in the index structure. When deleting or updates, the index should have virtually 100% used space.

However, a Reverse Key index will have the values reversed and dispersed evenly throughout the index structure. As index blocks fill, there will be a very remote chance of it being due to the maximum indexed value and 50-50 block splits will result. The PCT\_USED is likely therefore to be significantly less, averaging approximately 70-75% over time.

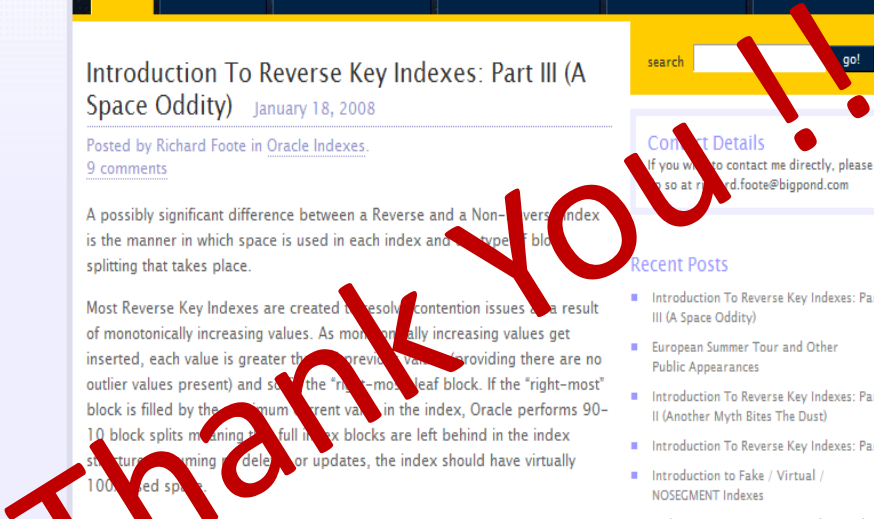
Therefore, for indexes with no deletions, a Reverse Key index is likely to be

search  go!

[Contact Details](#)  
If you want to contact me directly, please  
so at [richard.foote@bigpond.com](mailto:richard.foote@bigpond.com)

### Recent Posts

- Introduction To Reverse Key Indexes: Part III (A Space Oddity)
- European Summer Tour and Other Public Appearances
- Introduction To Reverse Key Indexes: Part II (Another Myth Bites The Dust)
- Introduction To Reverse Key Indexes: Part I
- Introduction to Fake / Virtual / NOSEGMENT Indexes
- 8 Things You May Not Know About Indexes
- 10,000 Hits Already !!
- Introduction To Linguistic Indexes - Part II
- DBMS\_STATS METHOD\_OPT default behaviour changed in 10g. Be careful ...
- Introduction To Linguistic Indexes - Part I
- Differences between Unique and Non-Unique Indexes



# Oracle Indexing Internals and Best Practices 5 Day Webinars



**8-12 October, 6-10 November 2018 (4 Hours Daily)**

*Of benefit to DBAs, Developers, Solution Architects and anyone else interested in designing, developing or maintaining high performance Oracle-based applications/databases.*

Examines most available index structures/options & discusses in considerable detail how indexes function, how/when they should be used & how they should be maintained. Also how indexes are costed & evaluated by the Cost Based Optimizer (CBO) & how appropriate data management practices are vital for an effective indexing strategy. Covers many useful tips and strategies to maximise the benefits of indexes on application/database performance & scalability.

Running between **10am-2pm Zurich Time 8<sup>th</sup> – 12<sup>th</sup> October 2018**

**[richardfooteconsulting.com/indexing-webinar/](http://richardfooteconsulting.com/indexing-webinar/)**

