

\*\*\* Create a basic little table

```
SQL> CREATE TABLE index_some_stuff (id number, status varchar2(20), description varchar2(50));
```

Table created.

\*\*\* Populate it with all STATUS values set to "PROCESSED"

```
SQL> INSERT INTO index_some_stuff SELECT rownum, 'PROCESSED', 'NOT REALLY INTERESTED WITH THIS ROW'
FROM DUAL CONNECT BY LEVEL <= 1000000;
```

Table created.

\*\*\* But set one row to have a STATUS of "BOWIE"

```
SQL> UPDATE index_some_stuff SET status = 'BOWIE', description = 'ROW OF INTEREST' WHERE id =
424242;
```

1 row updated.

```
SQL> COMMIT;
```

Commit complete.

\*\*\* So we have lots and lots of STATUS values that we're not really interested in and a "few" that we are ...

```
SQL> SELECT status, count(*) FROM index_some_stuff GROUP BY status;
```

STATUS	COUNT(*)
PROCESSED	999999
BOWIE	1

\*\*\* Let's create a normal index on the STATUS column

```
SQL> CREATE INDEX index_some_stuff_i_1 ON index_some_stuff(status);
```

Index created.

\*\*\* Let's collect stats on the table

```
SQL> exec dbms_stats.gather_table_stats(ownname=>'BOWIE', tabname=> 'INDEX_SOME_STUFF',
estimate_percent=> null, method_opt=>'FOR ALL COLUMNS SIZE 1', cascade=> true);
```

PL/SQL procedure successfully completed.

\*\*\* Let's search for those rows with a status of "BOWIE"

```
SQL> SELECT * FROM index_some_stuff WHERE status = 'BOWIE';
```

ID	STATUS	DESCRIPTION
424242	BOWIE	ROW OF INTEREST

1 row selected.

Execution Plan

Plan hash value: 290213492

Id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time
0	SELECT STATEMENT		500K	24M	1717 (2)	00:00:21
* 1	TABLE ACCESS FULL	INDEX_SOME_STUFF	500K	24M	1717 (2)	00:00:21

Predicate Information (identified by operation id):

1 - filter("STATUS"='BOWIE')

Statistics

```
-----  
0 recursive calls  
0 db block gets  
7765 consistent gets  
7754 physical reads  
0 redo size  
546 bytes sent via SQL*Net to client  
396 bytes received via SQL*Net from client  
2 SQL*Net roundtrips to/from client  
0 sorts (memory)  
0 sorts (disk)  
1 rows processed
```

\*\*\* add a histogram on the STATUS column

```
SQL> exec dbms_stats.gather_table_stats(ownname=>'BOWIE', tabname=> 'INDEX_SOME_STUFF',  
estimate_percent=> null, method_opt=>'FOR COLUMNS STATUS SIZE 5');
```

PL/SQL procedure successfully completed.

\*\*\* Let's search again for those rows with a status of "BOWIE"

```
SQL> SELECT * FROM index_some_stuff WHERE status = 'BOWIE';
```

```
-----  
ID STATUS DESCRIPTION  
-----  
424242 BOWIE ROW OF INTEREST
```

1 row selected.

Execution Plan

Plan hash value: 3846250962

```
-----  
| Id | Operation | Name | Rows | Bytes | Cost (%CPU) | Time |  
-----  
| 0 | SELECT STATEMENT | | 1 | 51 | 4 (0) | 00:00:01 |  
| 1 | TABLE ACCESS BY INDEX ROWID | INDEX_SOME_STUFF | 1 | 51 | 4 (0) | 00:00:01 |  
|* 2 | INDEX RANGE SCAN | INDEX_SOME_STUFF_I_1 | 1 | | 3 (0) | 00:00:01 |  
-----
```

Predicate Information (identified by operation id):

```
-----  
2 - access("STATUS"='BOWIE')
```

Statistics

```
-----  
0 recursive calls  
0 db block gets  
5 consistent gets  
0 physical reads  
0 redo size  
546 bytes sent via SQL*Net to client  
396 bytes received via SQL*Net from client  
2 SQL*Net roundtrips to/from client  
0 sorts (memory)  
0 sorts (disk)  
1 rows processed
```

\*\*\* And as expected, we have used our index.

\*\*\* But note that the index is predominately made up of "PROCESSED" values.

\*\*\* Perhaps we can improve on the 5 Consistent Reads we required for this index access ...

\*\*\* Let's create a function-based index that only stores the "BOWIE" STATUS values

```
SQL> CREATE INDEX index_some_stuff_i_2 ON index_some_stuff(DECODE(status, 'BOWIE', 'BOWIE', NULL))
COMPUTE STATISTICS;
```

Index created.

\*\*\* and create stats for the virtual column

```
SQL> exec dbms_stats.gather_table_stats(ownname=>'BOWIE', tabname=>'INDEX_SOME_STUFF',
estimate_percent=> null, method_opt=>'FOR ALL HIDDEN COLUMNS SIZE 1', cascade=> true);
```

PL/SQL procedure successfully completed.

\*\*\* Let's compare some stats on these two indexes

```
SQL> SELECT index_name, blevel, leaf_blocks, distinct_keys, num_rows FROM dba_indexes WHERE
owner='BOWIE' AND table_name='INDEX_SOME_STUFF';
```

INDEX_NAME	BLEVEL	LEAF_BLOCKS	DISTINCT_KEYS	NUM_ROWS
INDEX_SOME_STUFF_I_1	2	2924	2	1000000
INDEX_SOME_STUFF_I_2	0	1	1	1

\*\*\* Note that the "normal" index has a BLEVEL of 2 and 2924 leaf blocks

\*\*\* Note however the function-based index has a blevel of 0 and just the 1 leaf block ...

\*\*\* Let's run the same query now

```
SQL> SELECT * FROM bowie.index_some_stuff WHERE (DECODE(status, 'BOWIE', 'BOWIE', null)) = 'BOWIE';
```

1 row selected.

Execution Plan

Plan hash value: 2019748526

Id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time
0	SELECT STATEMENT		1	51	2 (0)	00:00:01
1	TABLE ACCESS BY INDEX ROWID	INDEX_SOME_STUFF	1	51	2 (0)	00:00:01
* 2	INDEX RANGE SCAN	INDEX_SOME_STUFF_I_2	1		1 (0)	00:00:01

Predicate Information (identified by operation id):

2 - access(DECODE("STATUS",'BOWIE','BOWIE',NULL)='BOWIE')

Statistics

```
0 recursive calls
0 db block gets
2 consistent gets
0 physical reads
0 redo size
546 bytes sent via SQL*Net to client
396 bytes received via SQL*Net from client
2 SQL*Net roundtrips to/from client
0 sorts (memory)
0 sorts (disk)
1 rows processed
```

\*\*\* and we note consistent reads have dropped from 5 down to just 2 ...