

*** First start with an Index Coalesce in a Non-ASSM tablespace

```
SQL> CREATE TABLE bowie_stuff (id NUMBER, name VARCHAR2(20));
```

Table created.

```
SQL> INSERT INTO bowie_stuff SELECT rownum, 'DAVID BOWIE' FROM dual CONNECT BY LEVEL <= 1000000;
```

1000000 rows created.

```
SQL> commit;
```

Commit complete.

*** Create index on ID column

```
SQL> CREATE INDEX bowie_stuff_i ON bowie_stuff(id);
```

Index created.

*** Delete 25% of rows

```
SQL> DELETE bowie_stuff WHERE mod(id,4) = 1;
```

250000 rows deleted.

```
SQL> commit;
```

Commit complete.

```
SQL> ANALYZE INDEX bowie_stuff_i VALIDATE STRUCTURE;
```

Index analyzed.

```
SQL> SELECT height, blocks, lf_blks, br_blks, btree_space, pct_used FROM index_stats;
```

HEIGHT	BLOCKS	LF_BKLS	BR_BKLS	BTREE_SPACE	PCT_USED
3	2304	2226	5	17848160	72

*** Note before and after value of redo size

```
SQL> SELECT n.name, s.value FROM v$mystat s, v$statname n WHERE s.statistic# = n.statistic# AND n.name = 'redo size';
```

NAME	VALUE
redo size	388379452

*** Coalesce the index

```
SQL> ALTER INDEX bowie_stuff_i COALESCE;
```

Index altered.

```
SQL> SELECT n.name, s.value FROM v$mystat s, v$statname n WHERE
s.statistic# = n.statistic# AND n.name = 'redo size';
```

NAME	VALUE
redo size (Diff 31,896,708)	420276160

*** Note that 31,896,708 bytes of redo generated

```
SQL> ANALYZE INDEX bowie_stuff_i VALIDATE STRUCTURE;
```

Index analyzed.

```
SQL> SELECT height, blocks, lf_blks, br_blks, btree_space, pct_used FROM
index_stats;
```

HEIGHT	BLOCKS	LF_BKLS	BR_BKLS	BTREE_SPACE	PCT_USED
3	2304	1671	5	13408160	90

*** Note that leaf blocks has reduced but not the number of blocks in the index segemnt ...

```
SQL> ALTER INDEX bowie_stuff_i SHRINK SPACE COMPACT;
ALTER INDEX bowie_stuff_i SHRINK SPACE COMPACT
*
```

```
ERROR at line 1:
ORA-10635: Invalid segment or tablespace type
```

*** Unable to use SHRINK as index is not in a ASSM tablespace ...

*** Repeat but this time create the index in an ASSM tablespace ...

```
SQL> drop table bowie_stuff;
```

Table dropped.

```
SQL> CREATE TABLE bowie_stuff (id NUMBER, name VARCHAR2(20));
```

Table created.

```
SQL> INSERT INTO bowie_stuff SELECT rownum, 'DAVID BOWIE' FROM dual CONNECT
BY LEVEL <= 1000000;
```

1000000 rows created.

```
SQL> commit;
```

Commit complete.

```
SQL> CREATE INDEX bowie_stuff_i ON bowie_stuff(id) TABLESPACE USERS;
```

Index created.

```
SQL> DELETE bowie_stuff WHERE mod(id,4) = 1;
```

250000 rows deleted.

```
SQL> commit;
```

Commit complete.

```
SQL> SELECT n.name, s.value FROM v$mystat s, v$statname n WHERE  
s.statistic# = n.statistic# AND n.name = 'redo size';
```

NAME	VALUE
redo size	580844640

*** Let's coalesce the index

```
SQL> ALTER INDEX bowie_stuff_i COALESCE;
```

Index altered.

```
SQL> SELECT n.name, s.value FROM v$mystat s, v$statname n WHERE  
s.statistic# = n.statistic# AND n.name = 'redo size';
```

NAME	VALUE
redo size (Diff 32,509,984)	613354624

*** Note redo similar to that of previous example ...

```
SQL> ANALYZE INDEX bowie_stuff_i VALIDATE STRUCTURE;
```

Index analyzed.

```
SQL> SELECT height, blocks, lf_blks, br_blks, btree_space, pct_used FROM  
index_stats;
```

HEIGHT	BLOCKS	LF_BKLS	BR_BKLS	BTREE_SPACE	PCT_USED
3	2304	1671	5	13408160	90

```
SQL> SELECT n.name, s.value FROM v$mystat s, v$statname n WHERE  
s.statistic# = n.statistic# AND n.name = 'redo size';
```

NAME	VALUE
redo size	613355428

*** Let's now shrink the index

```
SQL> ALTER INDEX bowie_stuff_i SHRINK SPACE;
```

Index altered.

```
SQL> SELECT n.name, s.value FROM v$mystat s, v$statname n WHERE  
s.statistic# = n.statistic# AND n.name = 'redo size';
```

NAME	VALUE
redo size (Diff 7,436,756)	620792184

*** Note that the shrink has generated an additional 7,436,756 bytes of redo

```
SQL> ANALYZE INDEX bowie_stuff_i VALIDATE STRUCTURE;
```

Index analyzed.

```
SQL> SELECT height, blocks, lf_blks, br_blks, btree_space, pct_used FROM  
index_stats;
```

HEIGHT	BLOCKS	LF_BKLS	BR_BKLS	BTREE_SPACE	PCT_USED
3	1720	1671	5	13408160	90

*** And it has also reduced the number of blocks allocated to the index segment (1720 down from 2304)

*** Now repeat with just SHRINK options

```
SQL> drop table bowie_stuff;
```

Table dropped.

```
SQL> CREATE TABLE bowie_stuff (id NUMBER, name VARCHAR2(20));
```

Table created.

```
SQL> INSERT INTO bowie_stuff SELECT rownum, 'DAVID BOWIE' FROM dual CONNECT  
BY LEVEL <= 1000000;
```

1000000 rows created.

```
SQL> commit;
```

Commit complete.

```
SQL> CREATE INDEX bowie_stuff_i ON bowie_stuff(id) TABLESPACE USERS;
```

Index created.

```
SQL> DELETE bowie_stuff WHERE mod(id,4) = 1;
```

250000 rows deleted.

```
SQL> commit;
```

Commit complete.

```
SQL> SELECT n.name, s.value FROM v$mystat s, v$statname n WHERE  
s.statistic# = n.statistic# AND n.name = 'redo size';
```

NAME	VALUE
redo size	781397216

*** Now use SHRINK SPACE COMPACT which defragments the index but does not actually deallocate any freed space

```
SQL> ALTER INDEX bowie_stuff_i SHRINK SPACE COMPACT;
```

Index altered.

```
SQL> SELECT n.name, s.value FROM v$mystat s, v$statname n WHERE  
s.statistic# = n.statistic# AND n.name = 'redo size';
```

NAME	VALUE
redo size (Diff 41,009,444)	822406660

*** Note it uses significantly more redo than the corresponding Index Coalesce
*** (roughly the cost of the Coalesce plus the cost of the following Shrink Space in previous example)

```
SQL> ANALYZE INDEX bowie_stuff_i VALIDATE STRUCTURE;
```

Index analyzed.

```
SQL> SELECT height, blocks, lf_blks, br_blks, btree_space, pct_used FROM  
index_stats;
```

HEIGHT	BLOCKS	LF_BKLS	BR_BKLS	BTREE_SPACE	PCT_USED
3	2304	1671	5	13408160	90

*** Note that the state of the index is identical to that of the Index Coalesce ...

```
SQL> SELECT n.name, s.value FROM v$mystat s, v$statname n WHERE
s.statistic# = n.statistic# AND n.name = 'redo size';
```

NAME	VALUE
redo size	822407348

*** Now let's complete the job and deallocate the freed space

```
SQL> ALTER INDEX bowie_stuff_i SHRINK SPACE;
```

Index altered.

```
SQL> SELECT n.name, s.value FROM v$mystat s, v$statname n WHERE
s.statistic# = n.statistic# AND n.name = 'redo size';
```

NAME	VALUE
redo size (Diff 167,692)	822575040

*** It required very little additional redo to complete the job

```
SQL> ANALYZE INDEX bowie_stuff_i VALIDATE STRUCTURE;
```

Index analyzed.

```
SQL> SELECT height, blocks, lf_blks, br_blks, btree_space, pct_used FROM
index_stats;
```

HEIGHT	BLOCKS	LF_BKLS	BR_BKLS	BTREE_SPACE	PCT_USED
3	1720	1671	5	13408160	90

*** And it has also reduced the index segment blocks down to 1720 ...