

*** Create another table loaded with a whole bunch of unique row values

```
SQL> CREATE TABLE case_search_2 as SELECT rownum id, 'PINK FLOYD' || rownum name FROM dual CONNECT BY level
<= 100000;
```

Table created.

*** Insert one special, easy to identify unique row ...

```
SQL> INSERT INTO case_search_2 VALUES (100001, 'THE WALL');
```

1 row created.

```
SQL> COMMIT;
```

Commit complete.

*** Create a normal, binary index on this new table

```
SQL> CREATE INDEX case_search_2_name_i ON case_search_2(name);
```

Index created.

*** Collect statistics on the table

```
SQL> exec dbms_stats.gather_table_stats(ownname=>'BOWIE', tabname=> 'CASE_SEARCH_2', cascade=> true,
estimate_percent=>null, method_opt=> 'FOR ALL COLUMNS SIZE 1');
```

PL/SQL procedure successfully completed.

*** Now set both NLS parameters required for Linguistic Indexes

```
SQL> ALTER SESSION SET NLS_SORT='GENERIC_M_CI';
```

Session altered.

```
SQL> ALTER SESSION SET NLS_COMP='LINGUISTIC';
```

Session altered.

*** Let's see what happens when we just want to select this one row ...

```
SQL> SELECT * FROM case_search_2 WHERE name = 'THE WALL';
```

```
      ID NAME
-----
100001 THE WALL
```

Execution Plan

Plan hash value: 1727594812

Id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time
0	SELECT STATEMENT		1	20	86 (6)	00:00:02
* 1	TABLE ACCESS FULL	CASE_SEARCH_2	1	20	86 (6)	00:00:02

Statistics

```
-----
1 recursive calls
0 db block gets
365 consistent gets
0 physical reads
0 redo size
470 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
```

*** We note that a FTS is performed even though there's an index on the name column and we just want the one row

*** Unfortunately, the existing index is a standard binary index and so is of no use and is ignored by the CBO with the NLS parameters set.

*** Let's perform a join between a table with an appropriate Linguistic Index (as created in the previous demo) with a table that only has a Binary Index

```
SQL> SELECT * FROM case_search cs1, case_search_2 cs2 WHERE cs1.name = 'ZIGGY' AND cs2.name = 'THE WALL';
```

ID NAME	ID NAME
100001 Ziggy	100001 THE WALL
100002 ZIGGY	100001 THE WALL
100003 ZiGgY	100001 THE WALL

Execution Plan

Plan hash value: 2400034535

Id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time
0	SELECT STATEMENT		1	43	89 (6)	00:00:02
1	MERGE JOIN CARTESIAN		1	43	89 (6)	00:00:02
* 2	TABLE ACCESS FULL	CASE_SEARCH_2	1	20	86 (6)	00:00:02
3	BUFFER SORT		1	23	3 (0)	00:00:01
4	TABLE ACCESS BY INDEX ROWID	CASE_SEARCH	1	23	3 (0)	00:00:01
* 5	INDEX RANGE SCAN	CASE_SEARCH_LING_NAME_I	1		2 (0)	00:00:01

Statistics

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

*** The Linguistic index on the CASE_SEARCH table is used however, the binary index is again ignored by the CBO and a FTS on the CASE_SEARCH_2 table is performed.

*** All binary indexes are of no use with nls_comp set to 'linguistic' and nls_sort not set to 'binary' ...