

---

## SQL queries slow in Aleph 18-up **\*\*MASTER record\*\***

- **Product:** Aleph
  - **Product Version:** 20, 21, 22, 23
  - **Relevant for Installation Type:** Dedicated-Direct, Direct, Local, Total Care
- 

### Description

The SQL shown below executed in 10 seconds in Aleph 16 but runs for more than 5 minutes in version 20.

```
SQL> select z13_title, z30_call_no from xxx50.z30, xxx50.z103, xxx01.z13 where  
substr(z30_rec_key,1,9) = substr(z103_rec_key,6,9)  
and z103_lkr_library = 'XXX01' and z103_lkr_doc_number = z13_rec_key and  
z30_item_process_status = 'GW' and z30_collection = 'PK';
```

After which it gets the following error:

```
ORA-01652: unable to extend temp segment by 128 in tablespace TEMP
```

There are:

4.3 million xxx01 z13 records, occupying 930 meg.

7.8 million xxx50 z103 records, occupying 540 meg.

6.1 million xxx50 z30 records, occupying 3.3 gig.

### Resolution

Oracle in the Aleph environment behaves as if it were using RBO (rule based optimizer) rather than the CBO (cost based optimizer). Since Aleph uses the RBO-like environment in its internal queries, the statistics, which make the CBO available, are disabled. Your (customer-written) queries may use an inefficient method of fetching data unless they are told explicitly to use RBO (**/\*+ ordered \*/**).

In this case, the query should look like this:

```
SQL> select /*+ ordered */z13_title,z30_call_no ... <etc.>
```

We suggest that all SQL queries involving the z103 be re-written to include the ORDERED or DYNAMIC SAMPLING hints in accordance with the principles described below. And that these hints be included in any other SQL which seems slower than it should be.

We tested the ORDERED, DYNAMIC SAMPLING, and INDEX hints with different queries on different databases.

We found that, for the multi-table queries I was investigating (see below), the INDEX hint (with single or multiple indexes) was always slower than one or both of the other hints and, therefore, have not included detailed analysis of that hint.

We find that, for the z30/z13/z103 query in the Description, the DYNAMIC SAMPLING hint

```
SQL> SELECT /*+ DYNAMIC_SAMPLING(2) ALL_ROWS */ z13_title, z30_call_no from xxx50.z30,  
xxx50.z103, xxx01.z13 where ...
```

gives a consistent result of about 40 seconds, regardless of the table order in the "from" clause. But the table order makes

a tremendous difference for ORDERED:

"from xxx50.z30, xxx50.z103, xxx01.z13" consistently completes in about 10 seconds\*;

"from xxx50.z103, xxx50.z30, xxx01.z13", about 50 seconds\*;

"from xxx50.z103, xxx01.z13, xxx50.z30", about 70 seconds\*;

the following three orders all take more than five minutes and/or give an "invalid number" error (see SKB 16384-26385, below):

from xxx50.z30, xxx01.z13, xxx50.z103 or

from xxx01.z13, xxx50.z103, xxx50.z30 or

from xxx01.z13, xxx50.z30, xxx50.z103 .

\* This is the first execution of the query. A second, third, etc., execution takes about 2/3 the time, presumably because of caching.

SQL> select count(\*) from z30 where z30\_item\_process\_status = 'GW' and z30\_collection = 'PK' ; retrieves 4,892 rows, whereas the complete z30, z13, and z103 tables all have millions of rows.

If I eliminate these z30-specific qualifiers ( z30\_item\_process\_status = 'GW' or z30\_collection = 'PK' ), the first query ("from duk50.z30, duk50.z103, duk01.z13", with the ORDERED hint) also gives the "invalid number" error.

Tests of other z30/z103/z13 queries on other servers / other databases gave similar results. We found cases where the ORDERED hint gives a worse result than no hint at all.

We did not see significant improvement from including \*both\* the ORDERED and DYNAMIC SAMPLING hints.

We believe that we can make the following three generalizations about multi-table SQL queries against Aleph tables:

1. For queries whose WHERE clause is retrieving relatively few rows from a particular table, including that table first in the FROM list, with the ORDERED hint, will almost always give the best result.

2. Queries involving ADM library tables, the ADM z103, and the bib z13 should have the following "FROM" table order when the ORDERED hint is used:

the ADM library table retrieving the fewest rows in the WHERE clause

the ADM library table retrieving the second fewest rows in the WHERE clause

other ADM tables

the ADM z103;

the BIB z13.

3. For queries whose WHERE clause is retrieving roughly equally large numbers of rows from \*all\* tables in the FROM list, the DYNAMIC SAMPLING hint may give a faster result than any of the ORDERED hints. Also, finding a good order for queries involving 5 or 6 tables with the ORDERED hint can be complicated; DYNAMIC SAMPLING doesn't require this optimization and almost always gives a good result. In view of the preceding, it seems that DYNAMIC SAMPLING is usually the simplest/best choice. You \*may\* want to include this in \*all\* of your SQL queries: we have found \*no\* case in which it slows a query down.

---

• **Article last edited:** 9-May-2016