
Parallel indexing, adding a new library and 2 servers

- **Article Type:** General
 - **Product:** Aleph
 - **Product Version:** 19.01
-

Description:

I am about to set up our system to do parallel indexing - we are having an external authority db walk, and will need to reindex after we reload the modified bib and authority records.

We will be adding more disk to the production server to do this indexing. So I will be opening a new library there. We are planning to do parallel indexing Option 2 (Oracle Import): after the new index tables are created, copy them into the original ts spots, and then drop the parallel indexing ones.

My question concerns our test server, to which we clone the production database from time to time. Should I be opening the new library there as well? Or would I just tell my Oracle DBA not to copy over the tablespaces used by the parallel indexes? (I am assuming it is possible to do that.)

Resolution:

If you did all the parallel indexing *between* clones to the test server, you could get away with not buying disk for the test server.

But you would still have to open up the new library on the test server, to match production -- it's just that there wouldn't be any data. I suggest copying the exact Aleph u-tree tables, with the parallel library definitions (in `aleph_start.private`, `library_relation`, `$alephe_root/xxxnn/`, etc.) from Prod to Test, but *not* having the big parallel tablespace you have on Prod on Test.

{The RMAN clone wants to have the exact tablespaces, with the exact sizes on Test as Prod, so you have two choices: (1) not do an actual clone: copy all the tablespaces except the parallel tablespace; or (2) temporarily drop the parallel tablespace at the time you do the RMAN clone.}

-
- **Article last edited:** 10/8/2013