
Multiple instances of p_manage_37?

- **Article Type:** General
 - **Product:** Aleph
 - **Product Version:** 18.01
-

Description:

I need to perform a fix on all of our bib records (delete 035 containing a specific string). In testing this in our test region I have found it takes slightly over 24 hours to process a 500,000 record range and we have almost 2 million. Is there any problem with me starting multiple manage_37 jobs from the command line, operating on separate record ranges? I'm thinking of running 8 concurrent jobs of 250,000 records each some weekend night.

I don't see any intermediate manage_37 files in \$data_scratch or \$alephe_scratch. Is this plan possible? Or is there a better way?

Resolution:

I'm uncertain about running multiple p_manage_37's -- on different record ranges -- at the same time.

It may be that this would work, but I don't actually have any experience with it. Some of the jobs write work files or sort-work files, and I there might potentially be conflicts, but perhaps not in the case of p_manage_37....

Assuming that the number of records with 035's containing this string is fewer than 10% of the database, then I think that a different approach is better:

First run p_ret_01 to locate the records containing an 035 with this particular string.

Then use that file as input to p_manage_37, so p_manage_37 will be reading *only* the records it needs to. I would *think* that this should be fast enough that a single p_manage_37 run over a weekend should be able to do all of them.

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