
ETL process: how long? what needs to be down?

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

Description:

We need to have clarification of what we should expect for the ETL specifically how long it might take, and what Aleph functionality will be unavailable during the process. For example, is it true that there should be no ARC activity while the ETL process is running?

Resolution:

I *think* the p_arc_01 is about 2-300,000 records per hour. So that would be 5 - 8 hours for your 1.5 million records.

There should be no ARC activity while the ETL is running.

Reading of the ALEPH OPAC and GUI could occur; updating should not. (Updating will result in the extract taking longer since Oracle wants to achieve a consistent image of the data. And could result in errors if the rollback segments are not big enough.) Since the Z69 OPAC statistics table is being updated by the OPAC, it is probably better for everything to be down. But you could try it with some things up; in the worst case, you would just need to re-run it.

Note: If there's a backlog of z07's waiting to be processed, you should stop ue_01 while p_arc_01 is running.

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