
How big should UNDOTBS1 be?

- **Article Type:** General
- **Product:** Aleph
- **Product Version:** 20, 21, 22, 22

Description:

The UNDOTBS1 was 95% full. I added another 1 Gb and now it went down to 68%. We did an indexing this past weekend, but nothing had changed other than that. [Note: You may see the message: ORA-30036: unable to extend segment by 8 in undo tablespace 'UNDOTBS1' .]

Right now the table is set to AUTOEXTENSIBLE=NO.

What's the best way to manage this table? Should I set AUTOEXTENSIBLE to yes?

Resolution:

The use of UNDO space is temporary (limited by the undo retention parameter). It is not the case that UNDOTBS1 grows indefinitely -- though it does grow to fill up the tablespace -- see the article "[UNDOTBS1 tablespace 99% full.](#)" (KB 16384-53011).

The distributed UNDOTBS1 is 2 gig. Looking with util o/13/1, we find that two large sites have 4.5 gig and 6 gig, respectively. Small-to-medium-sized sites should not need more than 4 gig.

Note: one site (with 2 million bib records) found that the Aleph2Alma urm Migration Tool (aleph_2_urm_invent_main_script) failed with the "unable to extend" error. Adding a datafile to increase the size of UNDOTBS1 to 4 gig. allowed the process to run successfully.

It's possible that the UNDO retention period could need to be increased from the distributed value of 10800, but this should not be necessary and should be done only if you receive an error indicating a problem with this value.

Setting this tablespace to AUTOEXTENSIBLE should not be necessary, but should do no harm as long as the UNDO retention period is not increased by a large amount.

The article "[UNDOTBS1 tablespace 99% full](#)" discusses the handling of expired segments in the UNDO tablespace.

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