



CDB Software, Inc.

CDB/RePart

May 5, 2005

Repartitioning Partitioned Tablespaces

Rebalancing partitions to avoid the 4 gig limit, or to simply enhance application performance is no easy task. Prior to CDB/RePart, the process was an eleven-step process:

1. Analyze the data to determine new limitkeys. (This is an iterative process)
2. Unload all the data.
3. Sort all the data.
4. Drop the object.
5. Recreate the object.
6. Load all the data.
7. Make a Copy.
8. Run Runstats against the object.
9. Rebind all plans.
10. Analyze access paths for performance impacts
11. Grant all authorizations.

Objects that are of substantial size as to need rebalancing make it hard to simply analyze the data to find new limitkey values. The Unload and SORT then become harder and more expensive as the amount of data grows. Whenever a Drop and Recreate is necessary the process comes with the risk of a recover.

With DB2 V6 (circa 1999-2000), IBM finally addressed this problem and delivered their solution. Starting with DB2 V6, IBM allows limitkeys to be altered with SQL without having to drop the table. But that is only half of the solution. IBM's whole solution is to run an SQL ALTER of the limitkeys and follow that with an IBM REORG. The reason to follow up the ALTER LIMITKEY with an IBM REORG is that the ALTER places the object in REORP (reorg pending) status. "REORP" status actually means the object

needs to be rebuilt, not reorganized. With IBM's solution, there are still many steps to the whole process that are not addressed and must still be managed and monitored by the DBA. There is no solution to the analysis phase, the Unload phase, the SORT phase and the Load phase. The latter three are simply rolled into an IBM REORG. And, the larger the data, the more time consuming the analysis, and the more expensive and difficult the IBM REORG.

In DB2 V8, IBM has introduced the REBALANCE keyword to the REORG. So, during a reorg, the rows will be rebalanced evenly across partitions. This is a very inflexible solution. There is no ability to limit the number of keys used, override the limitkeys, or even to know what the object will look like when it is finished.

The CDB Solution

1996, CDB introduced their repartitioning solution. It was a product called CDB/RePart. CDB/RePart was the first product of its kind and still remains the only total solution for managing partitioned tablespaces. In a single step, not eleven, CDB/RePart, with its simple syntax of REPART TABLESPACE A.B, will do the analysis, update the limitkeys and move the data with CDB's patented No-Sort technology. As part of this same process, a copy, rebinds and runstats may be done. Authorizations are not a concern as objects never get dropped and recreated.

The end result is a single step solution to achieve a fully rebalanced and fully reorganized tablespace. If for any reason something goes wrong, CDB/RePart is fully restartable. Simply resubmit the job and the process completes. As the data grows larger, no more intervention is required from the user.

CDB/RePart also goes well beyond simple rebalancing. CDB/RePart can also balance



partitions given a range of partitions, add or drop partitions, roll data down into prior partitions, archive data from partitions, and convert simple or segmented objects to partitioned tablespaces.

CDB/RePart gives the user the flexibility to limit the number of keycolumns used in the repartitioning as well as the ability to override the limitkeys. Most importantly, CDB/RePart gives the user the ability to see what the object will look like when the repartitioning is done. CDB/RePart can analyze the object while leaving it in RW and completely available to applications. The user can go through several iterations of the “What If” processing to determine the best fit of limitkeys, keycolumns, number of partitions, etc. for the business applications.

These and other features make CDB/RePart the only true repartitioning solution.

About CDB Software

CDB Software, Inc. is a leader in data management solutions for DB2 z/OS. CDB focuses its business on DB2 for z/OS to provide unique and innovative solutions that enable companies to expand their DB2 system to meet business needs while controlling the overall cost of the mainframe. Founded in 1985, CDB is a privately held corporation based in Houston, Texas with offices worldwide.

For more information visit:

www.cdbsoftware.com



CDB Software, Inc.
11200 Richmond Ave.
Houston, TX 77082