

APPENDIX H

Reorganizing the DB2 Catalog

IN THIS APPENDIX

- When Should the DB2 Catalog and Directory Be Reorganized?

Prior to DB2 V4, it was not possible to reorganize the table spaces in the DB2 Catalog and DB2 Directory because of the internal hash and link structures built into these databases. But for DB2 V4 and later releases, you can expediently reorganize the DB2 catalog and DB2 directory in a systematic manner using the REORG utility.

The DB2 catalog is the central repository for DB2 object and user metadata. DB2 is constantly referring to that metadata as it processes applications and queries. The physical condition of the table spaces and indexes that comprise the DB2 catalog is therefore a major component in overall DB2 subsystem performance.

Likewise, the DB2 directory contains internal control structures such as DBDs and skeleton cursor tables that can be accessed only by DB2 itself. The information in the DB2 directory is critical for database access, utility processing, plan and package execution, and logging. Efficient access to this information is quite critical.

Prior to DB2 V4, the only option for any type of “reorganization” activity was to run the RECOVER INDEX utility on DB2 catalog indexes. This rebuilt the indexes, but had no impact on the underlying data housed in the actual physical table space. Catalog reorganization is permitted on table spaces and indexes in the DB2 catalog database (DSNDB06) and on specific table spaces (SCT02, SPT01, and DBD01) in the DB2 directory database (DSNDB01). As of DB2 V6, of course, the REBUILD INDEX utility can be run to rebuild DB2 Catalog and Directory indexes.

When Should the DB2 Catalog and Directory Be Reorganized?

To determine when to reorganize the system catalog, DBAs can use most of the same basic indicators used to determine whether application table spaces should be reorganized. Although it always has been a wise course of action to execute RUNSTATS on the DB2 Catalog table spaces, it becomes even more important now that these table spaces can be reorganized. These statistics can be analyzed to determine when a REORG should be run. When RUNSTATS is run for a catalog table space, the statistics about that system catalog table space are gathered and then stored in the DB2 Catalog tables themselves.

In general, the following indicators and situations should be reviewed when determining when to reorganize your system catalog table spaces and indexes:

- An increase in the value of the near- and far-off position indicators (NEAROFFPOSF and FAROFFPOSF in SYSINDEXPART)
- An increase in the value of the near and far indirect reference indicators (NEARINDREF and FARINDREF in SYSTABLEPART)
- A decrease in cluster ratio (CLUSTERRATIOF in SYSINDEXES)
- An increase in leaf distance (LEAFDIST in SYSINDEXPART)
- When the DB2 Catalog and Directory data sets are not using a significant portion of their allocated disk space (PRIQTY)
- When the DB2 catalog and directory data sets contain a large number of secondary extents

For the SYSDBASE, SYSVIEWS, and SYSPLAN catalog table spaces, the value for the FAROFFPOSF and NEAROFFPOSF columns of SYSINDEXPART can be higher than for other table spaces before they need to be reorganized. Additionally, you can use REORG when it is necessary to move the DB2 Catalog and Directory to a different disk device.

Synchronizing System Catalog Reorganization

It is a more difficult prospect to determine when the DB2 Directory table spaces should be reorganized. The RUNSTATS utility does not maintain statistics for these “table spaces” like it can for the DB2 Catalog. However, it is possible to base the reorganization of the DB2 Directory table spaces on the reorganization schedule of the DB2 Catalog table spaces. In fact, in certain situations, it is imperative that specific DB2 Directory table spaces are reorganized when a “companion” DB2 Catalog table space is reorganized. The chart in Table H.1 provides information on keeping the DB2 catalog and DB2 directory table spaces “in sync.”

TABLE H.1 DB2 Directory Reorganization Indicators

When You REORG...	Be Sure to Also REORG...
DSNDB06.SYSDBASE	DSNDB01.DBD01
DSNDB06.SYSPKAGE	DSNDB01.SPT01
DSNDB06.SYSPLAN	DSNDB01.SCT02

These table spaces are logically related. DB2 requires that you reorganize them at the same time to keep them synchronized.

DB2 Catalog Reorganization Details

There are 20 DB2 Catalog table spaces and six DB2 Directory table spaces (refer to Tables H.2 and H.3). DB2 has different rules for different sets of these table spaces. There are three groupings of table spaces:

- Cannot be reorganized at all
- Can be reorganized using normal REORG procedures
- Can be reorganized using special REORG procedures

TABLE H.2 DB2 Catalog Table Spaces (DSNDB06)

Table Space	Definition
SYSCOPY	Contains image copy information (2 tables)
SYSDBASE	Contains database object information (13 tables)
SYSDBAUT	Contains database and database authority information (2 tables)
SYSDDF	Contains information about distributed DB2 connections (8 tables)
SYSGPAUT	Contains resource authority information (1 table)
SYSGROUP	Contains storage group information (2 tables)
V7 SYSGRTNS	Contains information about DB2 routines, such as functions and procedures (2 tables)
V7 SYSHIST	Contains historical statistics (8 tables)
V7 SYSJAUXA	LOB table space for Java JAR BLOB data (1 table)
V7 SYSJAUXB	LOB table space for Java source CLOB data (1 table)
V7 SYSJAVA	Contains information about Java programs (3 tables)
SYSOBJ	Contains object/relational and routine information (10 tables)
SYSPKAGE	Contains package and stored procedure information (8 tables)
SYSPLAN	Contains plan information (5 tables)
V7 SYSSEQ	Contains sequence information (pre-V8) (1 table)
V8 SYSSEQ2	Contains sequence information (V8) (2 tables)
SYSSTATS	Contains optimization statistics (6 tables)
SYSSTR	Contains translation and check constraint information (4 tables)
SYSUSER	Contains user authority information (1 table)
SYSVIEWS	Contains view information (4 tables)

TABLE H.3 DB2 Directory Table Spaces (DSNDB01)

Table Space	Definition
DBD01	Contains database descriptor information (one table)
SCT01	Contains skeleton cursor table information (one table)
SPT02	Contains skeleton package table information (one table)
SYSLGRNX	Contains recovery log range information (one table)
SYSUTILX	Contains utility processing information (one table)

There are only two table spaces in the first grouping of table spaces that cannot be reorganized at all: DSNDB01.SYSUTILX and DSNDB01.SYSLGRNX. Do not attempt to reorganize these table spaces as DB2 will not permit it.

The second grouping of table spaces must be processed differently than other table spaces:

- DSNDB06.SYSDBASE
- DSNDB06.SYSDBAUT
- DSNDB06.SYSGROUP
- DSNDB06.SYSPLAN
- DSNDB06.SYSVIEWS
- DSNDB01.DBD01

These six table spaces require special “handling and care.” Because they have a different internal configuration than most other table spaces, a different calculation is required for the size of the unload data set (SYSREC) used during the REORG utility. These table spaces contain internal links. Links are internal pointers that tie the information in their tables together hierarchically. A link can be thought of as a type of parent-child relationship in which, due to these links, the BUILD and SORT phases of the REORG utility are not executed.

The WORKDDN, SORTDATA, SORTDEVT, and SORTNUM options are ignored when reorganizing these table spaces. Also, the REORG utility cannot be restarted from the last checkpoint when used against these six table spaces. Instead, it must be restarted from the beginning of the phase. Finally, as mentioned before, a different set of steps must be executed during reorganization for these table spaces.

All other DB2 Catalog and DB2 Directory table spaces can be reorganized like any other DB2 table space. Keep in mind that the LOB table spaces in the DB2 Catalog are under the same restrictions as any other LOB table space regarding DB2 utilities.

Steps to REORG the Six “Special” Table Spaces

The following steps should be used when reorganizing the six “special” table spaces:

1. Calculate the size of the unload data set (SYSREC).

The SYSREC data set for the “special” table spaces has a different format than the other table spaces. This causes a special calculation to be required to determine its size. The equation to use is

$$\text{DATA SET SIZE IN BYTES} = (28 + \text{LONGROW}) * \text{NUMROWS}$$

NUMROWS is the number of rows to be contained in the data set and LONGROW is the length of the longest in the table space. For DSNDB06 table spaces, the value for LONGROW can be determined by running the following SQL statement:

```
SELECT MAX(RECLENGTH)
FROM   SYSIBM.SYSTABLES
WHERE  DBNAME = 'DSNDB06'
AND    TSNAME = 'name of table space to REORG'
AND    CREATOR = 'SYSIBM';
```

2. Ensure incompatible operations are not executing.
3. Start database DSNDB01 and DSNDB06 for read only access.
4. Run QUIESCE and DSN1CHKR utilities.
5. Take a full image copy of entire DB2 catalog and directory table spaces.
6. Start DSNDB01 and DSNDB06 for utility access.
7. Execute REORG utility.
8. Take a full image copy of entire DB2 catalog and directory table spaces.
9. Start table space and associated indexes for read/write access.

Furthermore, keep in mind that these six table spaces cannot be reorganized using specifying SHRLEVEL CHANGE. And finally, the SORTDATA, SORTDEVT, SORTNUM, and SORTKEYS options are ignored for these table spaces.

Steps to REORG Regular Table Spaces

The following steps should be used when reorganizing the remaining “regular” system catalog and directory table spaces:

1. Calculate the size of the unload data set (SYSREC) using the normal calculation:

$$\text{DATA SET SIZE IN BYTES} = \text{LONGROW} * \text{NUMROWS}$$

In this case it is unnecessary to add the additional 28 bytes to the length of the longest row. This is because these system catalog table spaces do not utilize links.

2. Ensure that incompatible operations are not concurrently executing (see the next section for an explanation of incompatible operations).
3. Start the table space and its associated indexes for read-only access.
4. Run CHECK INDEX on all indexes associated with the table space that is being reorganized.

5. Take a full image copy of the entire DB2 catalog and directory table spaces.
6. Start the table space and its associated indexes for utility access.
7. Execute the REORG utility.
8. Take a full image copy of the entire DB2 catalog and directory table spaces.
9. Start the table space and any associated indexes for read/write access.

These steps should be familiar to you because they closely follow the steps executed during the reorganization of an application data table space. There are several additional required steps added as precautions because of the critical nature of the DB2 catalog and directory.

NOTE

It is important to take a full image copy before and after reorganizing any DB2 catalog or directory table space.

Catalog Reorganization Restrictions

In addition to the procedures outlined previously, there are several restrictions on the manner in which the REORG TABLESPACE utility can be used with system catalog table spaces. First, recall that the SYSUTILX and SYSLGRNX table spaces in the DB2 Directory cannot be reorganized.

When reorganizing the DB2 Catalog (DSNDB06) and DB2 Directory (DSNDB01) table spaces, the following options cannot be used:

- The UNLOAD ONLY option is not permitted.
- Online REORG is not permitted for catalog and directory table spaces with links.
- The LOG YES option is not permitted as image copies are explicitly required following a catalog and/or directory reorganization.

Also, the reorganization of two specific table spaces are treated differently than any other in the manner in which they are tracked by DB2. Generally, DB2 records the reorganization of any table space in the SYSIBM.SYSCOPY system catalog table. However, DB2 records the reorganization of the DSNDB06.SYSCOPY and DSNDB01.DBD01 table spaces in the log instead.

You cannot collect inline statistics on the following DB2 Catalog and DB2 Directory table spaces:

DSNDB06.SYSDBASE	DSNDB06.SYSDBAUT
DSNDB06.SYSGROUP	DSNDB06.SYSPLAN
DSNDB06.SYSVIEWS	DSBDB06.SYSSTATS
DSNDB06.SYSHIST	DSNDB01.DBD01

Finally, in many 24x7 environments, it may be necessary to reorganize the system catalog and dictionary while it is being accessed. However, because of the central nature of the system catalog and directory to the operation of DB2, the following restrictions apply to concurrent activity during catalog reorganization:

- ALTER, DROP, and CREATE statements cannot be executed during the reorganization of any DB2 catalog or DB2 directory table space with the exception of SYSIBM.SYSSTR and SYSIBM.SYSCOPY.
- The BIND and FREE commands cannot be issued when the following table spaces are being reorganized: SYSIBM.SYSDBAUT, SYSIBM.SYSDBASE, SYSIBM.SYSGPAUT, SYSIBM.SYSPKAGE, SYSIBM.SYSPLAN, SYSIBM.SYSSTATS, SYSIBM.SYSUSER, and SYSIBM.SYSVIEWS.
- No DB2 utility can be running while SYSIBM.SYSCOPY, SYSIBM.SYSDBASE, SYSIBM.SYSDBAUT, SYSIBM.SYSSTATS, and/or SYSIBM.SYSUSER are being reorganized.
- No plan or package may be executed during the reorganization of SYSIBM.SYSPLAN and SYSIBM.SYSPKAGE.
- The GRANT and REVOKE statements cannot be issued when REORG is being run on SYSIBM.SYSDBASE, SYSIBM.SYSDBAUT, SYSIBM.SYSGPAUT, SYSIBM.SYSPKAGE, SYSIBM.SYSPLAN, and/or SYSIBM.SYSUSER.

The ability to reorganize the DB2 catalog and directory table spaces provides the DBA with a potent tool for his or her system tuning arsenal.

Take the Proper Image Copies

Finally, be sure to take a full image copy both before and after reorganizing any DB2 Catalog or DB2 Directory object. Failure to do so can result in an unrecoverable system table—a situation which is best avoided at all costs.