
CDB/Restore

CDB/Restore is a program that restores data from sequential files, (DSN1COPYs or Image Copies) or existing tablespaces, into a DB2 tablespace and rebuilds all associated indexes. CDB/Restore runs in a fraction of the time it takes the IBM DB2 RECOVER utility or the DSN1COPY service aid, followed by DB2 index rebuilds, to do the same function.

This chapter summarizes the primary features and capabilities of CDB/Restore.

Features

Automatic OBID Translate

CDB/Restore automatically translates object identifiers (OBIDs) for single-table tablespaces.

Update RUNSTATS

CDB/Restore can, at your option, update the DB2 catalog with the RUNSTATS values calculated at the end of the run. Restore will also produce a RUNSTATS report. The updating and report generation are independent facilities and can be used alone or together.

SYSCOPY Posting

CDB/Restore will post the SYSIBM.SYSCOPY Catalog Table. The run is posted as a LOG(NO) LOAD REPLACE. CDB/Restore can turn on the COPY PENDING status for the target tablespace as well.

Index RECOVER/REORG

All indexes on the tables in the tablespace will be completely rebuilt during restore processing. The products of CDB/Restore are a tablespace that can be accessed with SQL and fully reorganized indexes.

Different Inputs

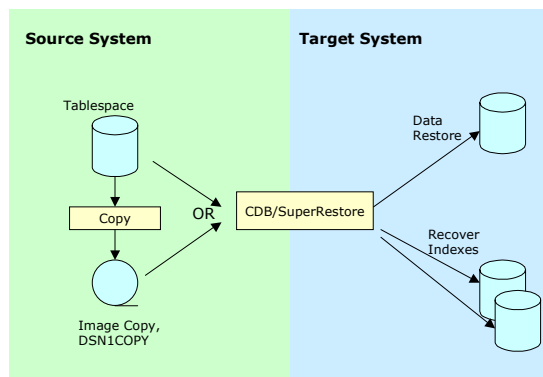
You may use either DSN1COPYs or DB2-compatible Image Copies as input to CDB/Restore. Basically, any sequential file conforming to the physical structure of a DB2 tablespace will work as an input dataset.

Index Independence

CDB/Restore will determine the requirements for building indexes at run time. This allows you to change the number of indexes or structure of an index very quickly and efficiently.

Theory of Operation

The input to CDB/Restore can be either a DB2 tablespace or a standard image copy taken from a tablespace. As shown in “*Restore Operation*”, CDB/Restore can be used to move data from a source system into a remote target system; on the other hand, the source and target systems can be the same, in which case CDB/Restore can be used to perform a “recover to copy” function.



Restore Operation

In the target system, CDB/Restore copies the input data into the target tablespace. At the same time, the indexes are recovered from the input data in parallel. This results in a process that is much faster than the sequential approach used by the standard DB2 utilities. This speed advantage increases with the number of indexes in the tablespace.

It should also be noted that the indexes in the source and target systems do not have to be the same. For example, if the target system is a decision support system with high rates of read access and little or no write access, the installation may wish to define a large number of indexes to optimized read access.

CDB/Restore Benchmarks

Restore (V7) Partitioned Tablespace, 110 Million Rows

