

Data Replication

Data Sheet



CDB Software, Inc.

CDB/[®]Auto-Restore V9.0.6

CDB/Auto-Restore for DB2[®] for z/OS[™] propagates data from one DB2 subsystem to another while simultaneously rebuilding indexes so that there is no subsequent step necessary to have your test or QA tables ready for processing. Through optimized multi-tasking and parallel processing, CDB/Auto-Restore efficiently moves data from VSAM to VSAM or from Image Copy to VSAM to create test data or restore dropped tables in a single step process without the need for a SYSXLAT table.

Key Benefits

- Built-In Automation
- No SYSXLAT table
- Optimized Parallel Processing/Multi-Tasking
- Single Step data propagation
- VSAM to VSAM or Image Copy to VSAM
- Indexes rebuilt at the same time as the table

The Challenge

You need to propagate data to your Test and QA tables. Whether you Unload from your production image copy then Load into QA or use DSN1COPY followed by Rebuild index, your still processing the data twice.

You need to restore a dropped object. You have an image copy dataset, DSN1COPY, or even the actual DB2 VSAM dataset of the dropped object. Again, DSN1COPY is cumbersome and inefficient as are Unload and Load

The CDB Solution


CDB/Auto-Restore is a unique product that restores your table from an Image Copy, DSN1COPY, or the actual DB2 VSAM dataset in one pass leaving you with Reorganized indexes and updated Runstats and will even optionally discard rows.

Built-In Automation

Only CDB has automation built into each utility. Once you realize what real automation is and work outside the boundaries that other products put on you, managing thousands of objects becomes easy and worry free. All dataset sizing, allocating, and extending are automated to your individual rules for error free execution. Error notification and job completion notification can be automatically built in to the process so that you are notified as soon as you restored tables are available.

Quick Availability


Whether you're propagating data or restoring a dropped table you'll have your data available fast with one simple job. No hunting for DBID, OBID, PSID combinations. While the data is being restored the indexes are also being rebuilt and all have Runstats collected. Your system is available and efficient as quickly as possible.



Efficient

CDB/Auto-Restore uses efficient I/O techniques directly to the DB2 VSAM datasets. By rebuilding indexes in parallel with the data propagation you get your data available in a fraction of the time and uses less CPU and I/O as well.

Data Sheet



*All benchmarks are based on a 60 Million Row Tablespace with 3 Indexes.

Important and unique features

• Runstats

CDB/Auto-Restore can collect statistics and update the DB2 Catalog during the restore. The target subsystem catalog is optionally updated during the Restore run.

• Discard/Purging

CDB/Auto-Restore can discard rows during the restore using a simple discard where clause with standard and powerful SQL.

• Automatic OBID Translation

If you're restoring the original table or propagating single table tablespaces you do not have to supply DB2's internal identifiers, they are translated automatically. If you are propagating multi-table tablespaces you map these identifiers just like you would in DSN1COPY, though, through built-in

automation, the translate map can be automatically generated.

• Multitasking/parallelism

You can Restore partitions in separate jobs to process them in parallel. CDB/Auto-Restore will keep work datasets for the NPIs in each job. When the last part is Restored the NPIs are automatically and quickly rebuilt using these datasets.

• DB2 Referential Integrity

You control whether Restore sets the Check Pending flag for tablespaces involved in DB2 RI. When you know you're restoring from a consistent state this will get you back online without running the lengthy Check Data utility.

• Outside of DB2

CDB Reorg operates on the DB2 VSAM datasets outside of DB2. CDB does not use DB2's Buffers or Logs. CDB uses a fraction of the CPU that IBM and the others use. Your applications are safe and unaffected.

• Built-in Automation

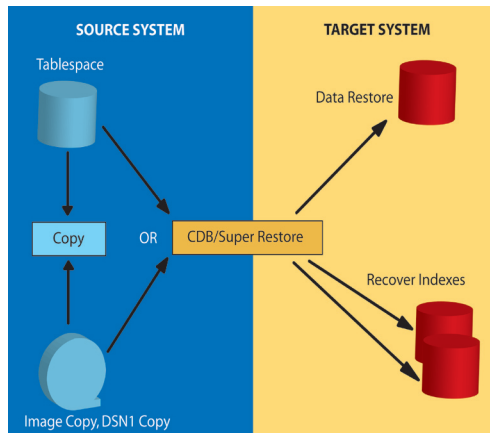
True automation is part of the utility, not a separate product that generates JCL. Automatic copy dataset selection to job completion notification can be tailored to fit your needs.

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