

Recovery Solutions

Data Sheet



CDB Software, Inc.

CDB/[®]Application Recovery V7.4.0

CDB/Application Recovery for DB2[®] for z/OS[™] is a high speed, single touch recovery system for groups of DB2 objects. Users may define groups of DB2 objects as a set called an “Application” using individual lists or wildcards. Log data is pre-processed for those defined objects on a regular basis and stored in compressed and ordered format. In the event of a need for a recovery of an application or an application object, a single statement such as RECOVER APPLICATION(SALES) will recover the whole application in a single pass to any point in time specified. Only objects that need to be recovered are recovered, while others are left in tact. Log is applied from the compressed ordered “REPOSITORY” eliminating the need for lengthy reads of the DB2 log. Your data is recovered quickly and efficiently without any extra work.

Key Benefits

- NO DATA CAPTURE
- Built-in Recover Manager
- Pre-processed log for optimal recovery times
- Reduced elapsed time through parallel processing
- Increased Parallelism through Server Technology
- Single statement to recover possibly thousands of objects

The Challenge

Since hardware has become very reliable, most recovery scenarios stem from application or operational errors. To recover from such errors you must perform a point-in-time recovery. Recovery often includes tables that are not directly affected but are referentially related to the affected table, often the whole application. Application recoveries should be set up when there is no pressure of a needed recovery but such recovery jobs are usually not maintained as applications change over time.

The CDB Solution

While other vendors try to make a single object recovery work for several objects, or possibly thousands of objects, CDB/ Application Recovery was designed for this purpose. CDB/ Application Recovery simplifies recovery. With one simple job, one simple statement recovers anything from one partition to the entire DB2 application.

All the objects needed for an application recovery are included using wildcards and objects are omitted using wildcards. As the application changes the wildcards include the changes.

To minimize any impact to the system, CDB/Application recovery does not require data capture to be enabled. This saves space and cycles due to reduction in DB2 logging.

Fast

CDB/Application Recovery uses advanced processing techniques exclusive to CDB. CDB server technology exploits your Zseries architecture to run the many small tasks in parallel using our efficient pre-processed log to reduce I/O. The DB2 logs are read once usually in advance of the application issue. The actual recovery is very efficient, fast, and doesn't tie up your DB2 log affecting other applications.

Safe

CDB/Application Recovery recovers your applications as defined in a calm thoughtful environment. At the time of recovery, there is no need to try to determine what objects should be recovered. There is no need to try to determine what objects belong to this application or which objects were affected by the problem. Simply submit a single statement job and your previously defined rules are followed. Objects that are not in need of recovery are automatically skipped.

Data Sheet

Important and unique features

- **Recovery Avoidance**

CDB/Application Recovery has a built in Recovery Manager. You can have CDB skip the recovery of any object that hasn't changed since the requested recovery point. This saves the restore from image copy and the forward recovery back to where it already is.

- **Asynchronous processing**

The processing of the DB2 logs is run in advance as Archive logs are written. The actual Application Recovery then runs very fast using input that has been efficiently consolidated. There is no need to read the DB2 log in time critical recovery periods.

- **Part of greater Architecture**

The asynchronous process of stripping the DB2 Log is part of a greater architecture. The data derived from stripping the log is also used by the CDB Delta, Audit, Clone, Copy, Application Recovery, and Active Standby products.

- **Server Technology**

CDB's Server Technology provides for maximum possible parallelism

by leveraging your Sysplex environment to distribute work across group members realizing the fastest possible elapsed time. You can control the throughput of your recovery based on need. If you need the recover to run even faster, use more servers. If system consumption is the concern, use less servers. Each recovery can be tailored to your needs.

- **Log Repository**

CDB's Striplog process preprocesses the DB2 log on a regular basis. Data that will be needed for recovery is kept in proprietary VSAM datasets in ordered format so that recovery can quickly access records it needs when it needs them.

The data in the repository is compressed and depending on the number of objects available for recovery, can usually be estimated to be 20-30% of the size of the full DB2 log.

- **Windows Console**

CDB Application Recovery is run with one simple job and CDB manages the rest. Many jobs are run on all nodes of your sysplex and the output is compiled in a simple Windows Console monitor.

You watch all processes on one simple GUI monitor.

- **Only Processes Completed UOWs**

CDB only processes completed UOWs so you can recover to any point in time without the need for a System Wide Point of Consistency or a system Quiesce. There is never any need for a time-consuming roll-back after a CDB recovery.

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