IBM PRESS

An Introduction to IMS

Your Complete Guide to IBM Information Management System

Barbara Klein, Diane Goff, John Butterweck, Kenny Blackman, Margaret Wilson, Moira McFadden Lanyi, Rick Long, Sandy Sherrill, and Steve Nathan

Second Edition

An Introduction to IMS

Second Edition

Database Recovery Control (DBRC) can be used to help ensure that there is always a recovery process available. Using DBRC to control database backup and recovery is not required but is highly recommended.

Online Programs and Recovery

IMS online transactions use dynamic backout to undo updates done in any incomplete unit of work. Online programs that abend are automatically backed out by the online system using the log records. In addition, if the system fails while an application program is active, any updates made by that program are automatically backed out when the system is restarted.

If the program is a batch message processing (batch message processing [BMP] program), the updates are automatically backed out to its most recent checkpoint. Because of this automatic backout, the recovery of individual databases is not needed.

At IMS restart time, if the emergency restart cannot complete the backout for any individual transactions, the databases affected by those updates are stopped, and DBRC is requested to set the recovery needed flag to ensure that a correct recovery is completed before the database is opened for more updates. In the case of dynamic backout failure, a batch backout or database recovery must be performed, depending on the reason for the backout failure.

DB Batch Update Programs and Recovery

DB batch update programs can use dynamic backout, as do BMP programs, if the following JCL changes are made:

- The `BKO=Y` parameter is set in the `EXEC` statement
- A DASD log data set is provided in the `IEFRDER DD` statement
- A Rollback (`ROLB`) call is issued in the program code for nonsystem abends

The dynamic backout backs out the updates to the last checkpoint in the log data set.

IMS Backup and Recovery Utilities

IMS provides the following utilities for recovering a database:

- **Database Image Copy utility**

  The Database Image Copy utility creates image copies of databases.

- **Database Image Copy 2 utility**

  The Database Image Copy 2 utility takes image copies of IMS databases by using the concurrent copy function of the Data Facility Storage Management Subsystem (DFSMS).

- **Online Database Image Copy utility**

  The Online Database Image Copy utility creates an as-is image copy of the database while it is being updated by the online system.

- **Database Change Accumulation utility**

  The Database Change Accumulation utility accumulates database changes from DL/I log tapes since the last complete image copy.

- **Log Archive utility**

  The Log Archive utility (DFSUARC0) produces a secondary log data set (SLDS) from a filled online log data set (OLDS) or a batch IMS SLDS. The utility runs as a z/OS batch job, and multiple log archive utility jobs can execute concurrently.

- **Database Recovery utility**

  The Database Recovery utility restores the database by using a prior database image copy and the accumulated changes from Data Language/Interface (DL/I) log tapes.

- **Batch Backout utility**

  The Batch Backout utility removes changes made to databases by a specific application program.

Another utility program, the Log Recovery utility (DFSULTR0), produces a usable log data set in the event of an operating system or hardware failure, thus enabling use of the log by the four principal programs of the recovery system.

For databases that consist of multiple data sets, recovery is done by individual data set. To recover a complete database that is composed of multiple data sets, database recovery must be performed for each of its component data sets.

Figure 11-2 illustrates the relationship between the backup and recovery utilities.

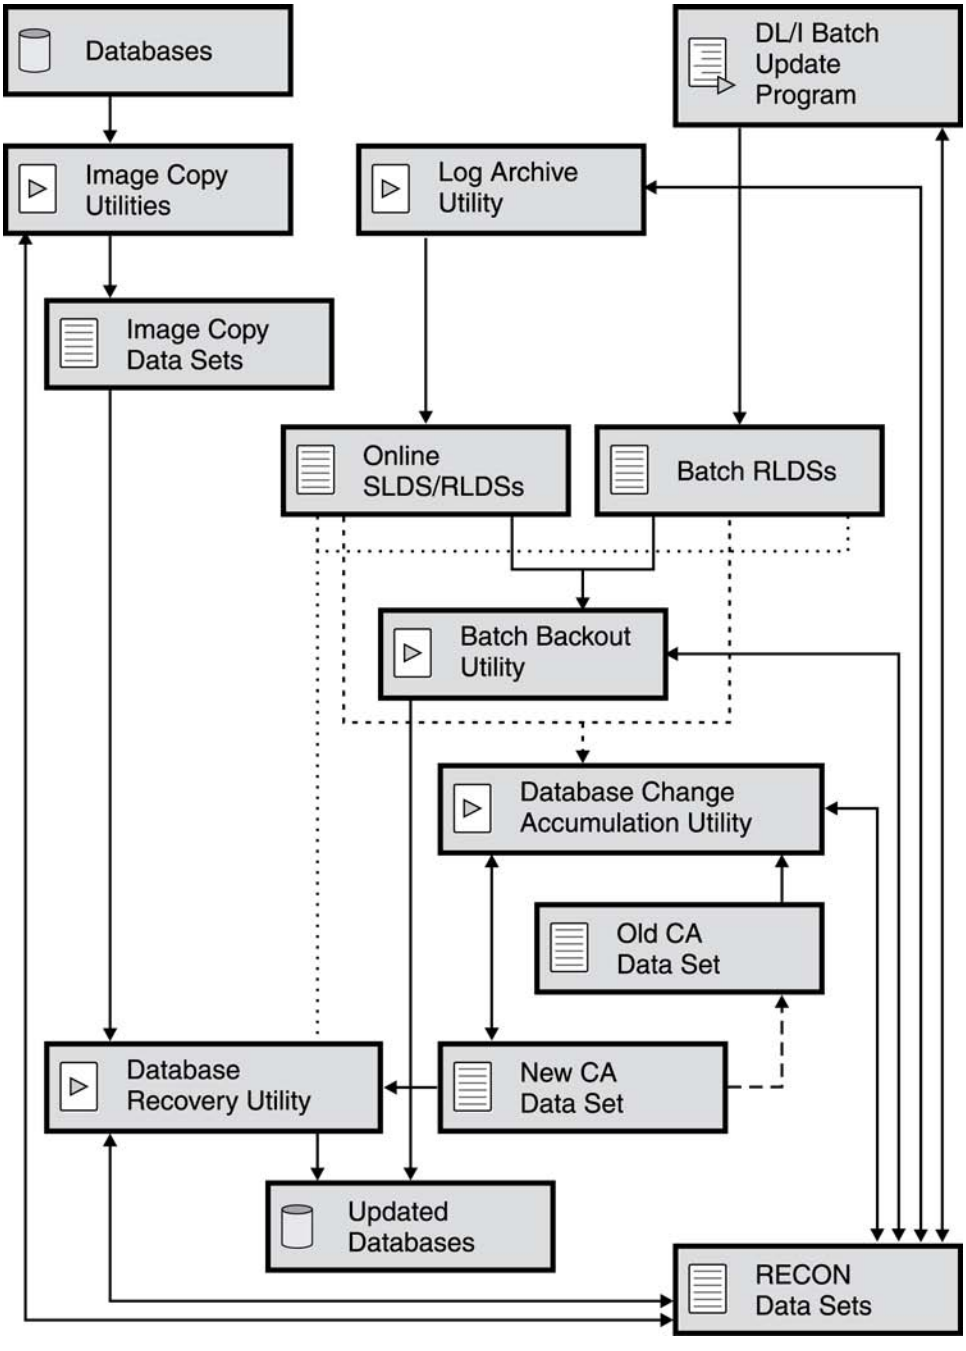

Figure 11-2 Overview of the Recovery Utilities

Database Image Copy Utility

The Database Image Copy utility (DFSUDMP0) creates a copy of each data set in the database. This copy is not an exact duplicate of the database data set, but it can be used to recover the database data set. The set of output data sets from the Database Image Copy utility is called an image copy. Each data set in the image copy is a sequential data set and can be used only as input to the Database Recovery utility (DFSURDB0). The Database Image Copy utility does not use DL/I to process the database, but uses track I/O instead. No internal checking occurs to determine whether all the IMS internal pointers are correct. There are tools available to run as part of the image copy process to do this checking. It is recommended that you periodically check the validity of the internal pointers.

There can be no changes to the database description (DBD) when a database is recovered using the Database Recovery utility. To make changes to the DBD, a database reorganization is needed to implement those changes.

Multiple databases and data sets can be copied with one run of the Database Image Copy utility. All data sets of a database should be copied at the same time: when there is no intervening database processing.

The Database Recovery Control (DBRC) facility can be used to generate the job control language (JCL) to run the Database Image Copy utility, if required.

A flow diagram of the Database Image Copy utility is shown in Figure 11-3.

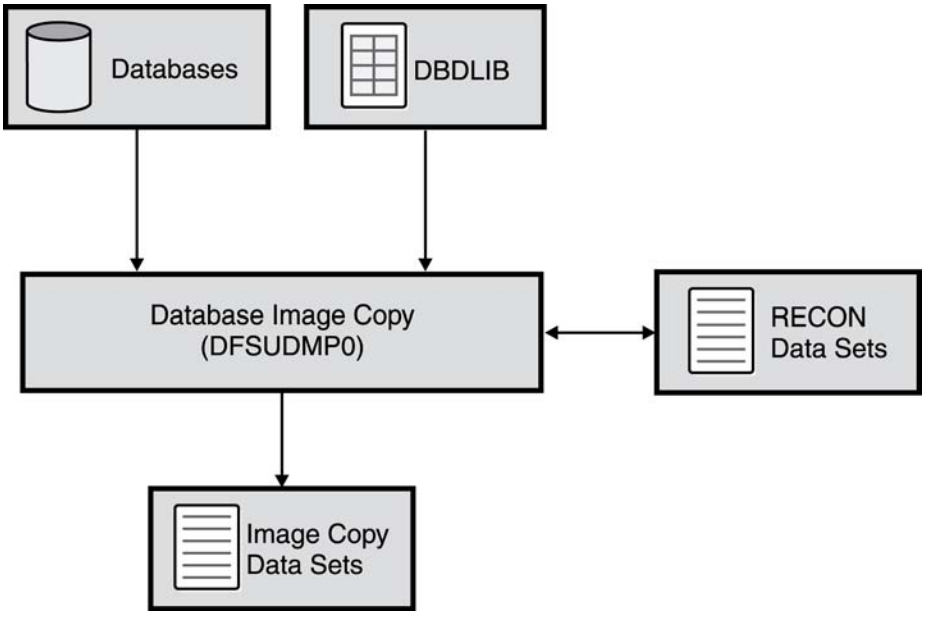

Figure 11-3 Input to and Output from the Database Image Copy Utility

Database Image Copy 2 Utility

The Database Image Copy 2 utility (DFSUDMT0) is similar to the Database Image Copy utility (DFSUDMP0). The Database Image Copy 2 utility has several advantages, however, in that it can take image copies while databases are unavailable for a very short time.

The DFSUDMT0 utility calls the concurrent copy function of DFSMS to make consistent image copies (that is, with no updates occurring while the utility is making the copy) or concurrent image copies (called *fuzzy* image copies) of an IMS database data set. The Database Image Copy utility (DFSUDMP0) cannot make concurrent image copies.

The concurrent copy function of DFSMS is a hardware and software solution that enables you to back up a database or any collection of data at a point in time and with minimum down time for the database. The database is unavailable only long enough for DFSMS to initialize a concurrent copy session for the data, which is a small fraction of the time that the complete backup process takes.

A flow diagram of the Database Image Copy 2 utility is shown in Figure 11-4.

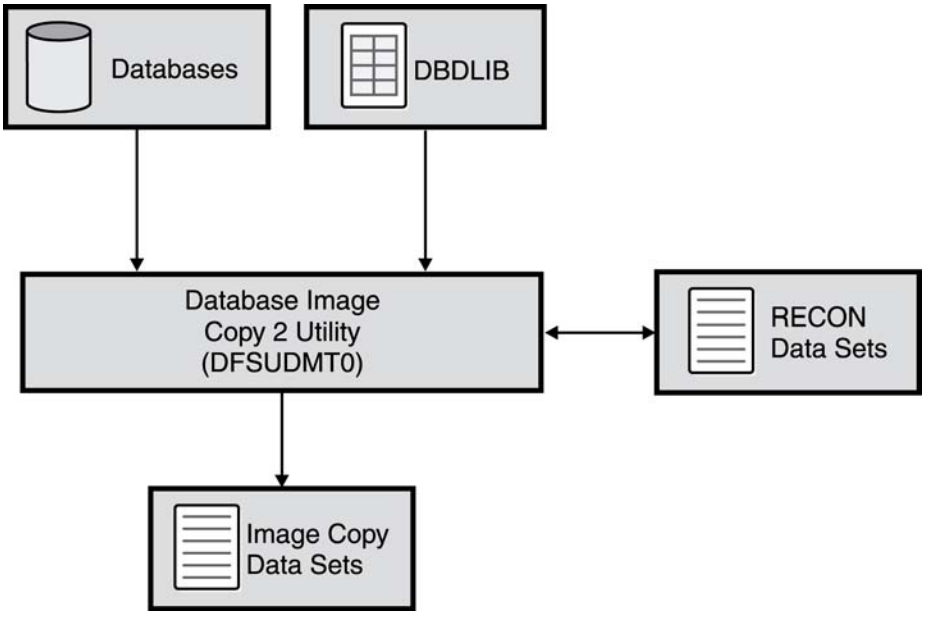

Figure 11-4 Input to and Output from the Database Image Copy 2 Utility

Online Database Image Copy Utility

The Online Database Image Copy utility (DFSUICP0) creates an as-is image copy of the database while it is being updated by the online system. The DFSUICP0 utility runs as a BMP program. You can use it only for hierarchical indexed sequential access method (HISAM), hierarchical indexed direct access method (HIDAM), and hierarchical direct access method (HDAM) databases. If IMS updates these databases while the utility is running, IMS requires

all logs for any subsequent recovery, including the log in use when you started this utility. IMS requires the logs because the image copy is not an image of the database at any particular point in time.

The output from the Online Database Image Copy utility is used as input to the Database Recovery utility.

A flow diagram of the Online Database Image Copy utility is shown in Figure 11-5.

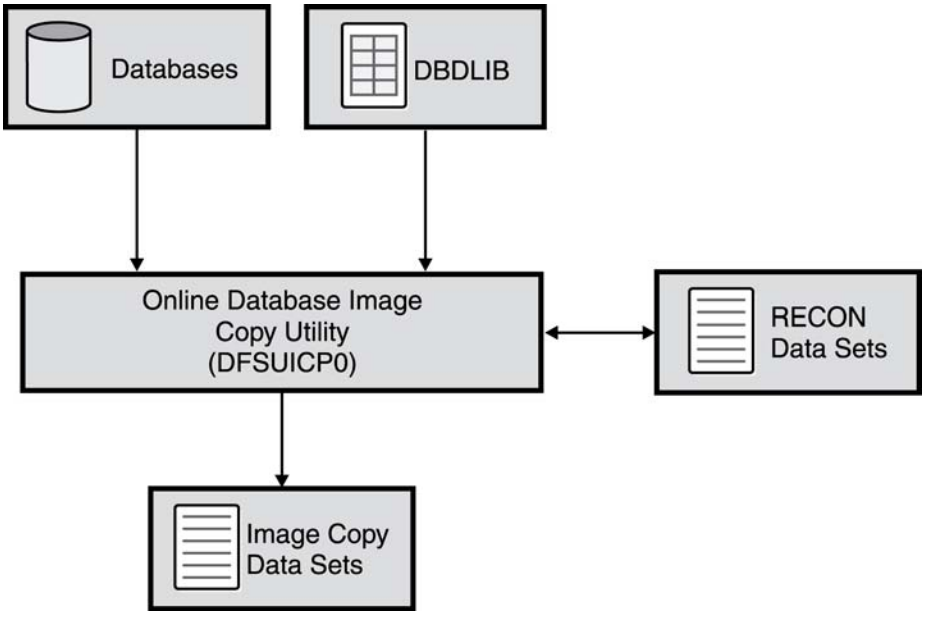

Figure 11-5 Input to and Output from the Online Database Image Copy Utility

Database Change Accumulation Utility

The Database Change Accumulation utility (DFSUCUM0) creates a sequential data set that contains only the database log records from all the log data sets that are necessary for recovery. This *change accumulation log data set* is used by the Database Recovery utility. The accumulation is done by sorting only the required log records in the physical record within data set sequence. Change accumulation provides efficient database recovery whenever needed. The number of log data sets that need to be kept is significantly reduced.

The Database Change Accumulation utility can be run independently of DL/I application programs.