

CONTENTS

1	Introduction.....	3
2	Automating the entire cloning process	3
3	Fast copy tools	3
3.1	Snaps tools.....	4
3.2	Mirroring tools.....	4
3.3	Point-in-time copy tools	5
4	Fast copy tools alone are not the answer.....	5
5	Primary considerations when cloning data	6
5.1	What version of DB2 is involved?	6
5.2	How much data needs to be cloned?	7
5.3	Where is the data located?	7
5.4	What kind of data access is needed?	7
6	Why are companies cloning DB2?	7
6.1	To increase productivity.....	7
6.2	What companies are saying	8
7	Conclusion.....	8

1 Introduction

Is 24x7 data availability truly possible?

Most companies think so and are creating DB2 clones to gain more use of their data. Did you know that you can have unlimited use of your existing DB2 data by replicating it and using it for other purposes?

Fast copy tools allow you to replicate a group of data within minutes to provide cloned environments onsite. However, until recently, these fast copy tools have been used almost exclusively for local or disaster recovery by replicating the data at the volume level, and then creating backups of the cloned data.

This is no longer the case. DB2 subsystem replicas are being used for testing, quality assurance or data warehousing environments, performance improvement, recovery initiatives, query-only access, and reporting.

As the demand for DB2 clones continues to increase, the job of the Database Administrator will become more complicated because of the cloned data access challenges. Providing DB2 subsystem clones in a timely fashion and from the same image is key.

This White Paper provides information on:

- automating DB2 subsystem data cloning
- the way that fast copy tools play a major role in the automation process
- cloning considerations
- the way that cloning automation software provides access to the cloned volumes from the same image in minutes versus hours or days with traditional utilities.

2 Automating the entire cloning process

To achieve the maximum amount of cloning automation, three main tools are needed:

1. Job scheduling software to automate the jobs.
2. Fast copy tools, such as snap, mirror, and point-in-time copy products to create the replica.
3. Cloning automation software to provide access to the replicated data from the same image (by renaming and cataloging it) and to update the cloned DB2 internals.

3 Fast copy tools

There are three basic categories of fast copy tools:

- SNAP tools (a combination of hardware and software), such as IBM FlashCopy, STK SnapShot, EMC TimeFinder/SNAP, and HDS Copy-on-Write Snapshot.

- Mirroring tools (a combination of hardware and software), such as EMC TimeFinder/Mirror, EMC SRDF and SRDF/A, HDS ShadowImage, HDS TrueCopy and IBM PPRC
- Point-in-time (PIT) copy tools (software only), such as Innovation Data Processing FDRPAS, Softek Replicator and Softek TDMF (with PIT enabled)

With these fast copy tools you can replicate entire DB2, DB2 PeopleSoft, or DB2 SAP subsystems in online (DB2 is suspended) or offline (DB2 is shut down) mode, to create usable clones.

3.1 Snaps tools

Snaps provide the appearance of an almost instantaneous copy of your data at a point in time. There are data set level and volume level snap tools. This White Paper will focus only on volume level snap tools.

For FlashCopy (NOCOPY) and SnapShot, the copy completes in a few seconds by constructing a bitmap to describe the source volume. The building of the bitmap is called the logical completion. The apparent copy time is mere seconds per volume. At the logical completion, the source volume and target volume contents are exact duplicates at that point in time.

The PIT copy is maintained and preserved by an on-demand copy of the data to the target volume when necessary. No data is physically copied from the source volume to the target volume except when data is about to be updated on the source volume or on the target volume.

The possible situations are:

- Read request to the source volume: The data is read from the source volume.
- Read request to the target volume: If the data was previously copied from the source volume to the target volume, the data is read from the target volume. If the data has not been copied to the target volume, the data is read from the source volume.
- Write request to the source volume: The data will be copied from the source volume to the target volume before the source volume update occurs. All updates are done in cache to prevent holding up any I/Os.
- Write request to the target volume: If the data was previously copied from the source volume to the target volume, the write request results in a normal write to the target volume. If the data has not been copied from the source volume to the target volume, the data will be copied to the target volume before the target volume update occurs. All updates are done in cache to prevent holding up any I/Os.

3.2 Mirroring tools

A mirror is caused by the creation and maintenance of a mirror relationship between the source and the target volumes. This relationship occurs onsite, as opposed to remotely. Each hardware vendor has different names or terms for the following processes but they are all very similar.

First, a relationship needs to be established between a source and target volume in plenty of time for the data to be copied to the target volume. At a designated point in time, a split or detach of the mirror relationship is initiated.

A split of a mirror guarantees that the contents of the target volume are an exact duplicate of the source volume at the point-in-time where the split occurs. The target volume is usually offline and needs to be re-labeled and varied online so the data can be used.

After the data is used, some tools allow the target data to be refreshed from the source without having to copy all of the data again – the data plays catch-up. Other tools require the mirror relationship to start all over.

3.3 Point-in-time copy tools

A point-in-time copy can be achieved using software-only copy tools such as Innovation Data Processing FDRPAS, Softek Replicator, and Softek TDMF (with PIT enabled). While FDRPAS and Softek TDMF are primarily thought of in terms of migrating data, they can also be used to replicate data.

A replica using these tools is caused by the creation of a paired relationship between the source and the target volumes. Source data is then copied to the target volumes. During the initial copy, updates to the source data are tracked. After the initial copy is complete, all tracked updates are applied asynchronously. This updating of the target happens over and over again until the desired time to create the PIT is reached. At a designated point in time, the data access to the source data has to be stopped so that the source and target volumes can be synchronized. After synchronization, the paired relationship is terminated.

4 Fast copy tools alone are not the answer

The challenge with copies is how to access the data once it has been created. Fast copy products do not manage the renaming, cataloging, and updating of the DB2 internals on the target volume data sets to produce a usable clone on the same image, nor do they update the target volume internals (VTOC, VTOCIX, and VVDS) to reflect the target volume names instead of the source volume names to enable the data to be used on the same image.

Some fast copy tools re-write the target VOLSER label with the source VOLSER name as an exact copy, while others retain the target VOLSER label. In all cases, the replica's data set and volume internal information, VTOC, VTOCIX, and VVDS, will be identical to the source. If you cannot access a non-sharing system, you cannot access the data!

The solution lies with cloning automation software like IBM DB2 Cloning Tool.

- DB2 Cloning Tool gives you access to data sets on target volumes by:
- Renaming and cataloging the data sets
- Fixing the target volume VTOC, VTOCIX, and VVDS conflicts to reflect the target volume names instead of the source volume names
- Updating all DB2 internal control information in the DB2 directory, BSDSs, and DB2 Catalog to enable access to the target DB2 subsystem.

5 Primary considerations when cloning data

5.1 What version of DB2 is involved?

The version of DB2 being used will directly affect how much downtime is required. Starting with DB2 V7, DB2 has commands to suspend (force any pending database writes to disk, suspend update activity, and flush the log buffers to disk) and resume (resume source DB2 activity) features that allow data to be cloned while it is online or up to reduce the impact of downtime.

When cloning an entire DB2 subsystem, the inability to easily update the DB2 internals (DB2 catalog, directory, and BSDSs) to reflect renamed data sets and target volumes can be extremely prohibitive without cloning automation software.

An online cloning solution often results in transactions in flight. These in-flight transactions, cloned to the target system, result in the same target system action that would happen on the source system if it were to have died at that same time and then restarted. The target restart is essentially an emergency restart of a failed system.

If you are cloning a DB2 subsystem online, you have to be absolutely certain that transactions in flight will not be backed out of the production DB2 subsystem when the target DB2 subsystem is brought up. Remember the target DB2 subsystem is a clone of the source DB2 subsystem including the DB2 internals that reflect the source names and volume serials. DB2 Cloning Tool for z/OS ensures that this does not happen.

To clone a DB2 subsystem, the following items must be performed in advance:

- Set up the target DB2 subsystem definitions so the target DB2 subsystem can be brought up. Procedures for setting up a DB2 subsystem can be found in the following documents: *DB2 UDB for OS/390 and z/OS V7 Administration Guide* and *DB2 UDB for z/OS V8 Administration Guide*.
- Set up a special maintenance mode ZPARM in advance to allow the DB2 catalog to be updated when the target DB2 subsystem is brought up in maintenance mode.
- Set up the DB2 Cloning Tool command, DB2SQL by using a Plan and Bind to automate the generation and issuance of all required DB2 SQL statements to update the DB2 catalog.
- If you are using DB2 V7, apply DB2 APAR PQ77546 (PTF UQ80188) to prevent the DB2 RTS function on the target DB2 system from corrupting tables on the source DB2 system. This is required for DB2 V7, but is automatically included in DB2 V8.

To create a successful DB2 subsystem clone, be sure to include all volumes in the clone that are related so that pieces of multi-volume data sets are not left behind.

5.2 How much data needs to be cloned?

The amount of data to clone will dictate which type of copy technology to use.

If thousands of volumes need to be cloned at once, the copy choice will involve using storage vendor copy tools. With these tools, the clone will be created on the same vendor DASD. Most often, however, the same type of data is not spread across DASD vendors so you will still be able to achieve your clone.

If hundreds of volumes need to be cloned at once, then the copy choices are extended to include using software vendor copy, which will clone the data on any vendor DASD.

5.3 Where is the data located?

Each vendor has rules about where the source volumes to clone and the target volumes as the recipient of the clone need to be located. Restrictions often dictate that the target volumes need to be in the same subsystem as the source volumes.

However, that does not mean that all source volumes have to be in the same subsystem. As long as there are an equal number of target volumes to pair with the source volumes in the same subsystem, the data can be cloned. Volumes in SMS storage groups can span subsystems.

Volume sizes have to be the same or larger and device types have to be the same.

5.4 What kind of data access is needed?

Recognize that the state of the replica will depend on the type of copy product used. Some copy tools will leave the replicated target volume online, while others will leave it offline. The target volume label may or may not be different than the source. Data set names will always be the same as the source including all of the volume internals, VTOC, VTOCIX, and VVDS. This knowledge will be crucial to understanding what steps need to be taken in order to use the clone.

6 Why are companies cloning DB2?

6.1 To increase productivity

By cloning an online system quickly, you can achieve virtually 24x7 availability, providing more services to your customers.

Online production downtime can be significantly reduced by using data cloning automation software.

6.2 What companies are saying

Based on a survey of DB2 Cloning Tool customers, here are some reasons why companies clone data:

- "The DB2 cloned target will be used by the SAP developers. They will begin changing the application for the next phase of code updates. They also develop a process to roll out the code to the production version of the application."
- "We clone DB2 instances to create TEST and production quality assurance environments for PeopleSoft HR and Financial. We use a production quality assurance environment for the User Community to test the changes that were made, and sign off that the changes can be rolled into the Production Environment."
- "We are using it to test out new functions and features of SAP."
- "We are moving a group of users to the cloned subsystem to lessen the performance impact on the production DB2."

7 Conclusion

Automating the data cloning process using DB2 Cloning Tool opens a whole new world of data access possibilities. DB2 subsystem clones can easily be created with DB2 Cloning Tool to allow unlimited database inquiry during batch processing, to create a truly effective test or acceptance test environment for applications to use or for data warehousing.

With a solution to access replicated data, new thought can be given to how your data can be accessed more frequently. The potential is immeasurable.



© Copyright Rocket Software Inc. 2006
© Copyright IBM Corporation 2006

IBM Corporation
Silicon Valley Laboratory
555 Bailey Avenue
San Jose, CA 95141
U.S.A.

Produced in the United States of America
11-06

All Rights Reserved
DB2, IBM, the IBM logo, the On Demand Business logo and
z/OS are trademarks of International Business Machines
Corporation in the United States, other countries or both.

Other company, product or service names may be trademarks
or service marks of others.

References in this publication to IBM products or services do
not imply that IBM intends to make them available in all
countries in which IBM operates.