


Copying PeopleSoft 8.1 DB2 Subsystem with Mainstar's Mirroring Solutions / Volume Conflict Rename using the FlashCopy Function of the IBM TotalStorage Enterprise Storage Server



PeopleSoft®

Written by: Edward A. Polley PeopleSoft Corporation and
Bill Worthington IBM Corporation

NOTICES AND DISCLAIMERS

Copyright © 2003 by International Business Machines Corporation and PeopleSoft Corporation.

No part of this document may be reproduced or transmitted in any form without written permission from IBM or PeopleSoft.

Product information and data has been reviewed for accuracy as of the date of initial publication. Product information and data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM and/or PeopleSoft may make improvements and/or changes in the product(s) and/or programs(s) described herein at any time without notice.

References in this document to IBM and PeopleSoft products, programs, or services does not imply that IBM or PeopleSoft intend to make such products, programs or services available in all countries in which either operates or does business.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR INFRINGEMENT. Neither IBM nor PeopleSoft shall have any responsibility to update this information. All products are warranted according to the terms and conditions of the agreements under which they are provided

The performance data contained herein was obtained in a controlled, isolated environment. Actual results that may be obtained in other operating environments may vary significantly. While each item had been reviewed for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere.

The responsibility for use of this information or the implementation of any of these techniques is a user responsibility and depends on the user's ability to evaluate and integrate them into their operating environment. Users attempting to adapt these techniques to their own environments do so at their own risk. IN NO EVENT SHALL IBM OR PEOPLESOFT BE LIABLE FOR ANY DAMAGE ARISING FROM THE USE OF THIS INFORMATION, INCLUDING BUT NOT LIMITED TO, LOSS OF DATA, BUSINESS INTERRUPTION, LOSS OF PROFIT OR LOSS OF OPPORTUNITY.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM or PeopleSoft patents or copyrights.

Trademarks

The following are trademarks, or registered trademarks of PeopleSoft, Inc in the United States, other countries, or both – PeopleSoft, PSOFT, PS HRMS, and PS North American Payroll.

The following are trademarks, or registered trademarks of the IBM Corporation in the United States, other countries, or both –DB2, Enterprise Storage Server, FlashCopy, IBM, MVS, TotalStorage, and z/OS.

The following are trademarks, or registered trademarks of the Mainstar Software Corporation in the United States, other countries, or both – Mainstar and Mirroring Solutions.

TEAM MEMBERS.....	5
PROJECT OVERVIEW	6
ENVIRONMENTAL SETUP.....	6
ENVIRONMENTAL SETUP.....	7
Hardware.....	7
DB2 Volume layout.....	8
Software	8
COPY PROCESS	10
COPY METRICS	15
FLASHCOPY ONLY	15
CLONED DB2 SUBSYSTEM	16
IMAGE COPY METRICS	17
RMF FINDINGS	18

TEAM MEMBERS

This project combined management and technical skills to push forward new concepts requiring a teamwork approach to combining these management and technical skills. We were fortunate on this project to have the support of the following:

- Eva Lau, WW Alliance Executive, IBM Storage Systems Group
- Eric Gray, Director Performance and Benchmarks PeopleSoft
- Humair Ghauri, Manager PeopleSoft
- Doug Sallee, Project Coordinator PeopleSoft
- Bob Perego, Mainstar

There were a number of skilled technicians who made the project successful they were:

- Elizabeth McGhee, Mainstar
- Bill Worthington and Charlie Burger, IBM
- Edward Polley and John Gray, PeopleSoft.

We wish to acknowledge the support of the PeopleSoft Operations Group who played a key role in setting up the z/OS environment and running critical jobs; Jim Ghinazzi, Jon Nolting and Baldeva Prasad.

There were many others who played supporting roles but important roles in this project, thanks to all.

PROJECT OVERVIEW

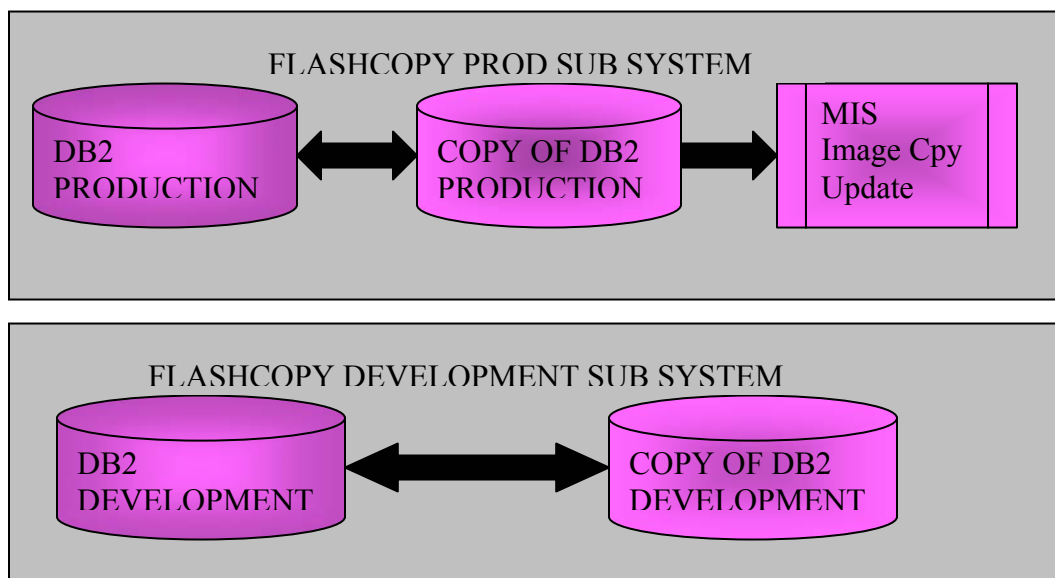
Business has a critical need for quick, accurate DB2 backup to support both their production copy/recovery and their development copy/restoration requirements. Prior to this paper, the process for recovery and restoration of entire DB2 subsystems often required the better part of a day. Moreover, there needed to be a target LPAR for the new subsystem distinct from the LPAR of the source subsystem.

In the fourth quarter of 2002 PeopleSoft, Mainstar and IBM formed a joint project team to test the practicality of quickly copying an entire PeopleSoft DB2 Subsystem within the same source subsystem LPAR. The cloned system was used to offload the original subsystem. The tests were conducted on an IBM TotalStorage Enterprise Storage Server (codenamed “Shark”)(“ESS”) and an eServer z900 server, using the ESS’s FlashCopy advanced function and Mainstar’s Mirroring Solutions / Volume Conflict Rename (MS/VCR) software.

This paper documents the hardware, software and processes needed to quickly backup and initiate processing on a different DB2 subsystem

Entire production DB2 subsystems can be copied in mere minutes to a new standalone DB2 subsystem. This new subsystem offloads MIS functions, traditional Image Copies and even Update activity (synchronization not discussed herein) helping to free the production system to devote more time to business critical jobs. This process is extremely quick and far more efficient than traditional Image Copy Restores.

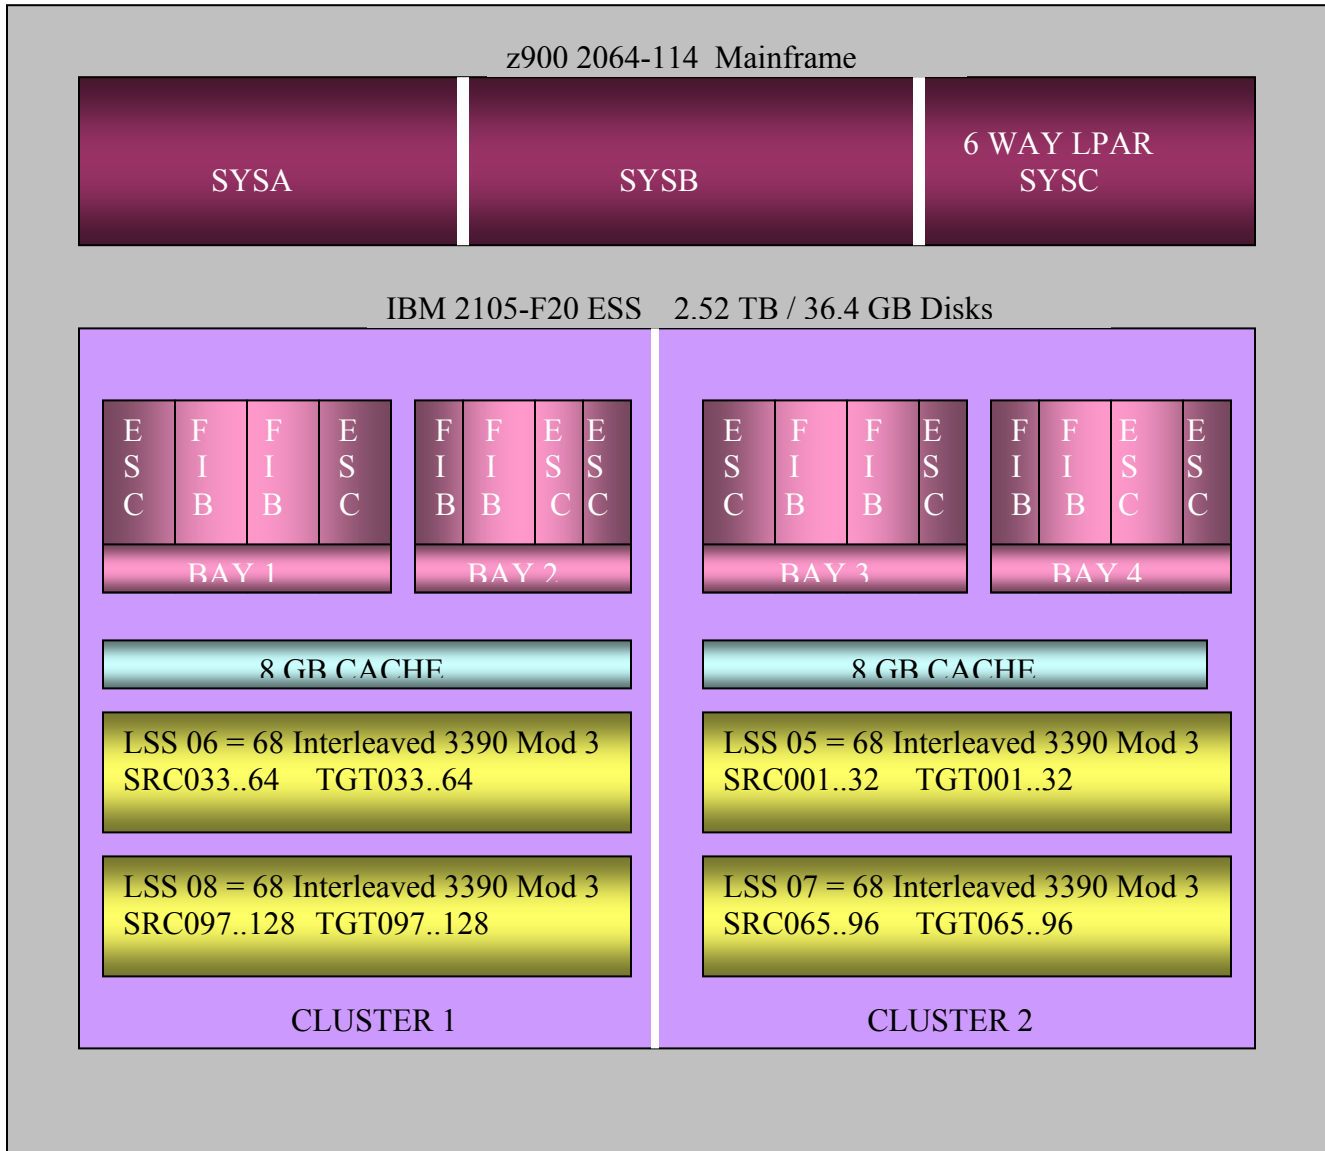
In addition, both production and developmental DB2 subsystems can be copied to ‘offline’ volumes as static data and used to quickly restore the original subsystem.



ENVIRONMENTAL SETUP

Hardware

The following chart shows the overall hardware configuration. We had access to the 6-way LPAR known as SYSC and an ESS Model F20.



There were 128 source disk volumes labeled ‘SRC*’ and of course 128 target volumes labeled ‘TGT*’. It is a requirement that the source and target volumes of FlashCopy be in the same ESS Logical Subsystem (LSS).

DB2 Volume layout

We tested with a full production-sized PeopleSoft North American Payroll database consisting of 9 databases, 78 tablespaces, 8054 tables and 14,828 indexes. This is important for the Mainstar software because the greater the number of database objects the more work must be completed to bring up the cloned DB2 subsystem. It also intended to reflect what customers may be running in their production environments.

The fully loaded large size benchmark database, (90,000 employees) was spread across the 128 source volumes without DBA intervention. The DB2 Stogroup included all 128 ‘source volumes.’

We normally isolate our DB2 catalog, directory, BSDSs, and LOGs on separate non-application data volumes for security reasons, so we also copied 5 additional volumes (SP4A40/1, SP4B40/1, SP4C40) as source to (SP4A42/3, SP4B42/3, SP4C41) as target.

In all we copied 133 volumes of data.

Software and Advanced Functions

- i. PeopleSoft 8.1g North American Payroll
- ii. DB2 for OS/390 Version 7.1
- iii. ESS FlashCopy.
- iv. Mainstar’s MS/VCR Version 1 Release 2G

FlashCopy Description

FlashCopy is a feature of the ESS. It provides an instant or point-in-time (PiT) copy of a logical volume giving an instantaneous copy or view of what the original data looked like at a specific point-in-time – sometimes called a time zero (T₀) copy. FlashCopy allows accessing both the source and target volumes even before the physical copy process has completed. By creating an instant image of the source, FlashCopy enables applications using either the source or target copy to operate with only the minimal interruption while the FlashCopy relationship is being established.

A FlashCopy relationship is a set of “extents,” where an extent is the unit of storage that the ESS copies. This relationship exists from the time the FlashCopy is initiated until the ESS finishes its background task of copying all data from the source extents to the target extents. (There is a Persistent FlashCopy option which enforces the relationship until it is explicitly withdrawn.) Invoking FlashCopy with DFSMSdss allows the target volume to retain its volume label or to have a new label specified. Our approach changes the label so that the targets can remain accessible in the SYSC LPAR.

Mainstar’s MS/VCR

Mainstar’s Mirroring Solutions / Volume Conflict Rename (MS/VCR) software offering gives access to “cloned” datasets on FlashCopied volumes. These clones have an inherent limitation. Although the target volume label reflects either the source or one specified for the target, the internal data – VTOC, VTOCIX, VVDS and dataset names all reflect the source volume label. MS/VCR is designed to rename and catalog the cloned data, fix the VTOC, VTOCIX and VVDS.

Additionally, MS/VCR updates DB2’s internal control information. This is done using the DB2UPDATE component which conditions the cloned DB2 subsystem by modifying the directory, BSDSs, and catalog.

COPY PROCESS

In this section, we detail the 9 easy steps which are required to quickly copy an entire DB2 subsystem as a new subsystem. It is recommended that the reader have a working knowledge of JCL and DB2. The steps are described in the sequence they must be executed.

- 1. BCSCLEAN**
- 2. STOP SOURCE DB2**
- 3. FLASHCOPY**
- 4. RESTART SOURCE DB2**
- 5. COPYCHECK**
- 6. RENAME**
- 7. DB2UPDATE**
- 8. DSNJU003**
- 9. START TARGET DB2**

The following text is the actual JCL used on this project. It is expected that the reader with appropriate modifications could use this JCL to obtain similar results.

1. **BCSCLEAN** – This job is not needed on the first run, but only if the copy step is retried. The critical DDs are JOURNAL and SYSIN. JOURNAL points to the Mainstar Journal dataset which the Mainstar job BCSCLEAN uses to complete its cleanup. Sample JCL follows:

```
//EPBCLEAN JOB , 'VCR BCSCLEAN', CLASS=Z, MSGCLASS=X
//STEP1 EXEC PGM=VCR00010
//STEPLIB DD DSN=SYS5.MAINSTAR.V1R200G.LOAD, DISP=SHR
//VCRINI DD DSN=EP.VCR.PARMLIB(VCRINI), DISP=SHR
//BCSRECS DD DSN=EP.VCR.WRK.BCSRECS, DISP=OLD
//JOURNAL DD DSN=EP.VCR.JRNL, DISP=OLD
//VCRPRINT DD SYSOUT=*
//SYSIN DD *
        BCSCLEAN
        JOURNAL(JOURNAL)
```

2. **STOP SOURCE DB2** – use DB2 command to stop source DB2 subsystem, this stop is very brief, about 3 minutes. It is highly recommended to stop DB2. Otherwise the target subsystem may have in-flight threads and it will not come up smoothly. Note that we only stopped the Payroll databases because there were no other active databases during FlashCopy, but if there were other active Databases we recommend they all be stopped.

-STOP DB (PAYROLL)

3. **FLASHCOPY** – This job copies source volumes to target volumes extremely quickly, usually in several minutes, see timing section below. The critical DDs are JOURNAL and SYSIN. SYSIN identifies the maximum volumes to copy in a task (COPYCMDLIMIT) and the number of tasks (DSSTASKS), the copy process is so quick there was little advantage to altering these parameters. SYSIN also identifies the source volumes and target volumes, notice below that clause FROM-VOLSER uses a mask of SRC* to copy from all volumes with a VOLSER beginning with SRC. SRC* covers 128 volumes, the remaining 5 volumes are individually listed starting with SP4A40.

The TO-VOLSER clause identifies which volumes received the data, notice the use of the TGT* mask. The clause USERCATALOGS identifies how to copy the DB2 catalog datasets.

Since this job creates the important journal file, it is best to have the first step an IDCAMS delete of the journal file.

Once this job is done, the source database can be restarted.

```

//EPCP24X4 JOB ,'VCR COPY',CLASS=Z,MSGCLASS=X
//STEP0 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
  DEL EP.VCR.JRNL
  DEL EP.VCR.WRK.UCATBKUP.*
  SET MAXCC=0
//STEP1 EXEC PGM=VCR00010
//STEPLIB DD DSN=SYS5.MAINSTAR.V1R200G.LOAD,DISP=SHR
//VCRINI DD DSN=EP.VCR.PARMLIB(VCRINI),DISP=SHR
//BCSRECS DD DSN=EP.VCR.WRK.BCSRECS,DISP=OLD
//JOURNAL DD DSN=EP.VCR.JRNL,DISP=(,CATALG,),
// RECORG=KS,KEYLEN=64,KEYOFF=0,UNIT=SYSDA,
// LRECL=600,SPACE=(CYL,(10,10))
//VCRPRINT DD SYSOUT=*
//SYSIN DD
COPY
  DATA-MOVER (COPYCMDLIMIT(24)
                DSSTASKS(4))
FROM-VOLSER (SRC*
              SP4A40,SP4A41,SP4B40,SP4B41,SP4C40)
TO-VOLSER (TGT*
           SP4A42,SP4A43,SP4B42,SP4B43,SP4C31)
USERCATALOGS(
              CATALOG.DB1M CATALOG.DB1X
              CATALOG.DS1M CATALOG.DS1X )
JOURNAL-DDN(JOURNAL)

```

Note it is recommended that the DSSTASKS parameter equal the number of LSSs containing source data. We recorded no significant benefit when we increased the DSSTASKS parameter greater than four.

4. **RESTART SOURCE DB2** – As soon as the brief Copy job completes the source DB2 can be brought back online. As indicated in the timeline below this is just several minutes.

```
-START DB(PAYROLL)
```

5. **COPYCHECK** – *This job is optional.* It monitors the progress of the ESS background copy. It is important to note that once the job completes (see step 3) the source database is available.

```

//EPCOPYCK JOB ,'VCR COPYCHECK',CLASS=Z,MSGCLASS=X
//STEP1 EXEC PGM=VCR00010

```

```

//STEPLIB DD DSN=SYS5.MAINSTAR.V1R200G.LOAD,DISP=SHR
//VCRINI DD DSN=EP.VCR.PARMLIB(VCRINI),DISP=SHR
//JOURNAL DD DSN=EP.VCR.JRNL,DISP=(,CATALG,),DISP=SHR
//VCRPRINT DD SYSOUT=*
//SYSIN DD
COPYCHECK
WAIT(12)
JOURNAL-DDN(JOURNAL)

```

6. **RENAME** – Critical job for the target DB2 subsystem as this one renames the tablespace High Level Qualifier (HLQ). Notice that the ‘rename masks’ are source HLQ DB1M to target HLQ DB1X and source subsystem name DS1M to target subsystem name DS1X.

```

//EPRNAME JOB , 'VCR COPY', CLASS=Z, MSGCLASS=X
//STEP0 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
DEL EP.VCR.WRK.BCSRECS
DEL EP.VCR.WRK.IDCAMS
DEL EP.VCR.WRK.VOLBKUP
/*
//STEP1 EXEC PGM=VCR00010
//STEPLIB DD DSN=SYS5.MAINSTAR.V1R200G.LOAD,DISP=SHR
//VCRINI DD DSN=EP.VCR.PARMLIB(VCRINI),DISP=SHR
//SORTWRK01 DD UNIT=SYSDA,SPACE=(CYL(10,10))
//SORTWRK02 DD UNIT=SYSDA,SPACE=(CYL(10,10))
//SORTWRK03 DD UNIT=SYSDA,SPACE=(CYL(10,10))
//SORTWRK04 DD UNIT=SYSDA,SPACE=(CYL(10,10))
//SORTWRK05 DD UNIT=SYSDA,SPACE=(CYL(10,10))
//SORTWRK06 DD UNIT=SYSDA,SPACE=(CYL(10,10))
//VCRPRINT DD SYSOUT=*
//DRSTATS DD SYSOUT=*
//MSPRINT DD SYSOUT=*
//JOURNAL DD DSN=EP.VCR.JRNL,DISP=OLD
//BCSRECS DD DSN=EP.VCR.WRK.BCSRECS,DISP=(,CATLG,)
//MSWORK DD DSN=EP.VCR.WRK.IDCAMS,DISP=(,CATLG,)
//VOLBKUP DD DSN=EP.VCR.WRK.VOLBKUP,DISP=(,CATLG,)
//SYSIN DD *
RENAME
SPEED
MAX-TASKS(9)
JOURNAL-DDN(JOURNAL)
RECATALOG(N)
DATACLASS(SOURCE)

```

```

MGMTCLAS(SOURCE)
STORCLAS(SOURCE)
RENAME-MASKS (
    DB1M.*.*      DB1X.*.*
    DS1M.*.*      DS1M.*.*

```

7. **DB2UPDATE** – This critical job updates the BSDS on the target subsystem. Note the source and target DB2 pairing on the DB2UPDATE clause.

```

//EPDB2UPD JOB ,'VCR DB2 UPDATE',CLASS=Z,MSGCLASS=X
//STEP1 EXEC PGM=VCR00010
//STEPLIB DD DSN=SYS5.MAINSTAR.V1R200G.LOAD,DISP=SHR
//VCRINI DD DSN=EP.VCR.PARMLIB(VCRINI),DISP=SHR
//BSDS01 DD DSN=DS1X.BSDS01, DISP=SHR
//BSDS02 DD DSN=DS1X.BSDS02,DISP=SHR
//DBD01 DD DISP=OLD
//          DSN=DS1X.DSNDBC.DSNDB01.DBD01.I0001.A001
//SYSIN DD *
DB2UPDATE
    DB2-HLQS(DS1M DS1X  DB1M DB1M)
    JOURNAL-DDN(JOURNAL)

```

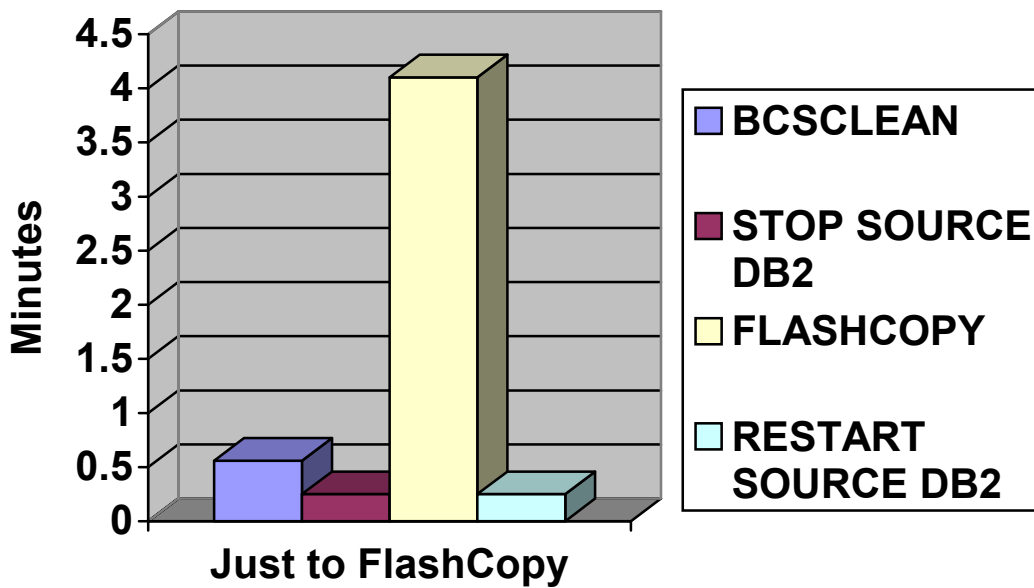
8. **DSNJU003** – This optional job is needed if you are using DDF, it updates the BSDSs with the new communication information. The critical DD is SYSIN, its values vary for each installation here are the values we used.
DDF
LOCATION=DB2DS1X.LUNAME=DB2APL1X,PORT=5118,RESPORT=5119,PA
SSWORD=

9. **START TARGET DB2 SUBSYSTEM.**

COPY METRICS

FLASHCOPY ONLY

The following section shows the run times of the three steps needed to FlashCopy 133 volumes of data and bring up the source DB2 subsystem. Notice that the FLASHCOPY step is in essence completed for 133 volumes in just 4 minutes and the source DB2 subsystem is down for just 5 minutes. This is remarkable considering the amount of data copied.



CLONED DB2 SUBSYSTEM

The following chart shows the total run time to not only copy the data but to bring up a cloned DB2 subsystem in the same LPAR. This totals to 36.6 minutes and shows all the tasks needed to bring up an entirely new DB2 subsystem. Note that the COPYCHECK time is optional, we choose to wait until the background copies completed in order to isolate the background copy RMF data. The RENAME job can be started as soon as the FlashCopy job is complete. This reduces the total run time to **just 9.6 minutes**


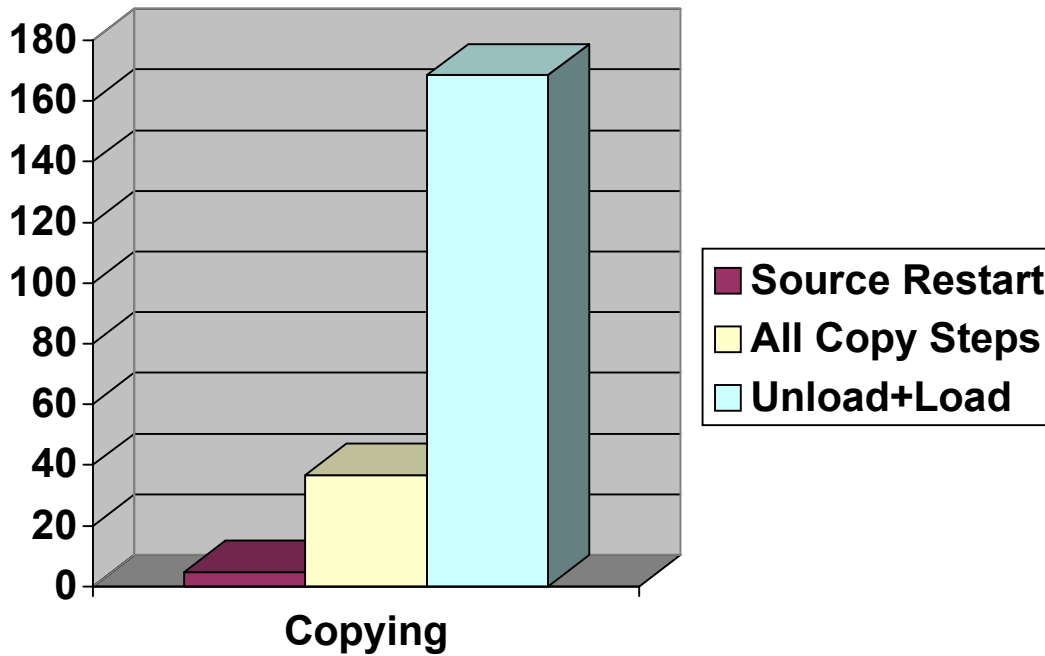


IMAGE COPY METRICS

One alternative to FLASHCOPY with MS/VCR is to create a shell DB2 subsystem and then to use DB2 UNLOAD and LOAD utilities to populate the data. For this task we would have to copy 128 volumes of data to tape (UNLOAD) then read it back into (LOAD) our new subsystem.

Since FlashCopy and Image Copy only copy the data in each VTOC or tablespace respective there is no architectural advantage to Image Copy. We used a PeopleSoft Production sized Payroll system at 50 GB of data. The jobs were run concurrently with a maximum of 8 at any one time as this was the most our processor would support. The following chart compares Unload/Load times compared to FlashCopy. Note that we used the new UNLOAD utility available in Version 7.1 DB2.

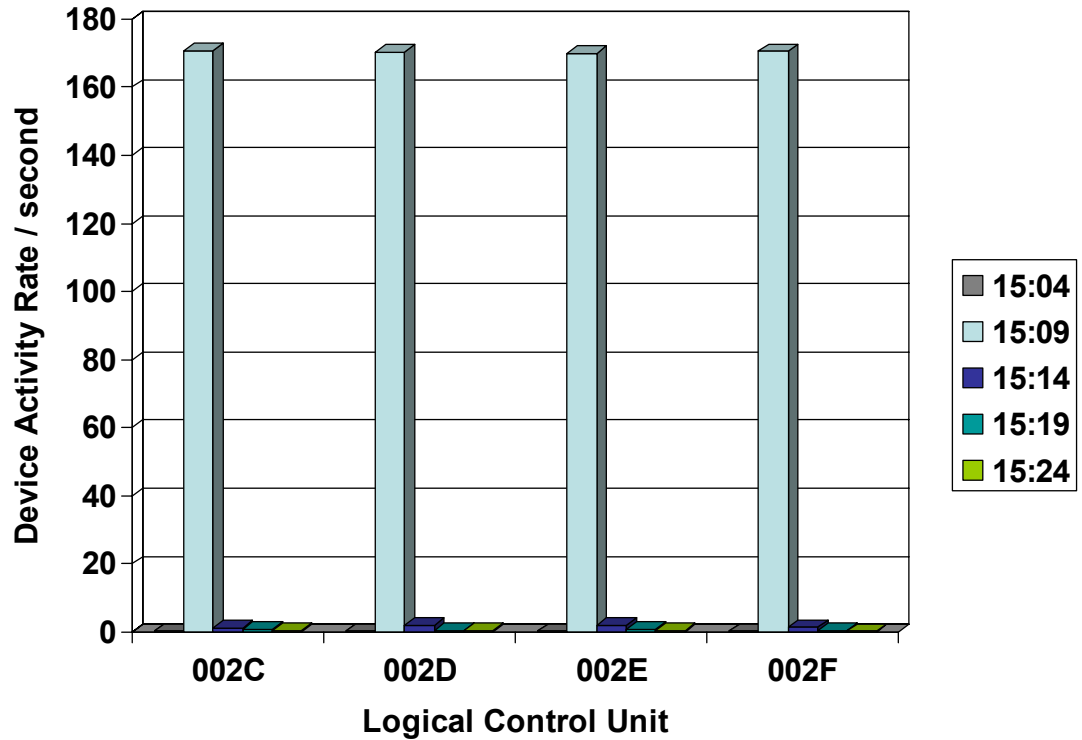


Performance Findings

The following charts show the FlashCopy results as collected by RMF.

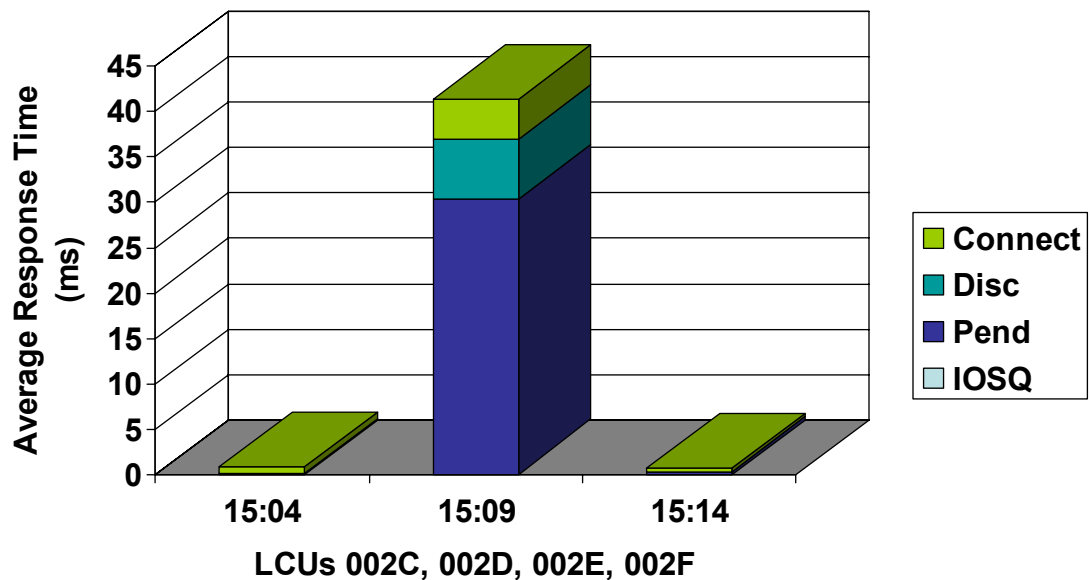
The first chart shows the burst of I/O activity as FlashCopy was initiated. Notice that there is little or no activity before or after the 15:09 interval, which is the time that FlashCopy is being invoked for the DB2 database and its associated tablespaces, tables and indices.

I/O Activity



The second chart shows that there was also an associated increase in response time while Mainstar's MS/VCR established the FlashCopy relationship between the source and target volumes for the payroll database. The "pend" time is high because, during FlashCopy initialization, the ESS reads the VTOC information of all the volumes being copied at once creating the burst of I/O activity in the previous chart. This briefly saturates the channels causing a 'channel busy' condition which is reflected as "pend" time.

Combined LCU Response Time



Once the FlashCopy relationship was established, the ESS began to physically copy the data from each source volume to the target volume as a background task without having to wait for all the FlashCopy relationships to be established. Note that both the source volumes and target volumes were available for use during this time. The next chart shows that the copying was started during the time that the FlashCopy step was being run and continued for approximately the next fifteen minutes. (The optional CopyCheck job step confirmed this.)

FlashCopy Background Copy Tracks Copied per Second

