HP 3000

Practical Guide To Mirrored Disk/iX

Commercial Systems Division
November 1996

Authored by: Roy Bittinger

1.0 INTRODUCTION

1.1 OVERVIEW

HP 3000 customers today require continuous access to their application data. Solutions which provide high availability of data despite a variety of planned and unplanned interruptions are in high demand. Mirrored Disk/iX is a software product designed to provide high availability for disks and I/O cards on MPE/iX systems. By automatically maintaining a copy of all data on two separate disks (called partners) at all times, Mirrored Disk/iX eliminates application downtime due to disk failure. This replication is transparent to the user and application. While Hewlett-Packard is a recognized leader in disk drive reliability, even the most reliable disk drives can fail. Prudent system managers invest in proven technology such as Mirrored Disk/iX to shield against unexpected disk failures.

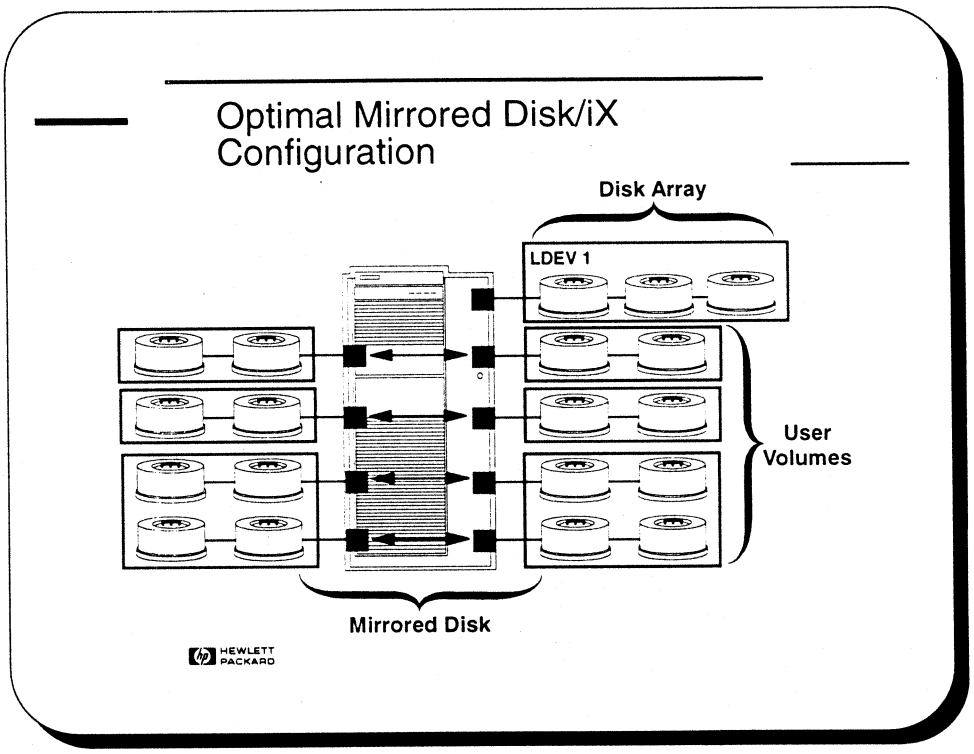

1.1.1 Key Benefits of Mirrored Disk/iX

High Data Availability

If a disk fails, users and applications are unaffected. The system availability and applications continue to function without interruption.

On-line Backup

Users can backup their data using TurboSTORE/iX Online Backup. Mirrored Disk/iX does not affect TurboSTORE/iX, nor does TurboSTORE/iX affect Mirrored Disk/iX. In the past before online versions of TurboSTORE/iX were available, customers performed Split-volume backups which utilized the split volume feature of Mirrored Disk/iX. This feature allowed users to take one of the pairs of the partner volumes offline so a backup could be performed on the offline pair, while the application continued running on the other copy of the data. As soon as the backup was complete, the volume would be brought back online and the pairs would resynchronize. This technique, split-volume backup, is no longer

recommended. Because Mirrored Disk/iX is designed to increase availability, performing a split-volume backup which compromises availability is contradictory to the use of the solution.

Recovery from Disk Failure

Should a disk failure occur, the system will automatically detect the drive failure. The Mirrored Disk/iX software discontinues mirroring at the same time and continues to run applications using the data on the "good" disk. This ensures that the application continues without interruption in the event of a disk failure.

Resume Mirroring

After mounting a new disk drive, assuming the disk required either repair or replacement, Mirrored Disk/iX synchronizes the new drive with the data from the surviving partner. Mirrored Disk/iX will ensure that the partner disk then becomes identical to its pair and regular mirroring is resumed.

Data Consistency

The system writes to both disk partners of a mirrored pair at the same time to ensure that data is consistent between them.

1.1.2 Reason For This Document

The purpose of this practical guide is to assist you in understanding what Mirrored Disk/iX is all about. Sometimes gaining a bit of information on how exactly the product works helps to make it clear whether the solution is the right product for your environment.

The guide is not intended to be used instead of the Mirrored DIsk/iX User's Guide. This guide should be used in conjunction with the User's Guide.

This document is fairly technical in nature and was written with the Operations Manager in mind. Many of the steps in here help the "implementor" of Mirrored Disk/iX understand exactly what it takes to set up the environment and achieve higher levels of data availability.

Should you have any questions during the reading of this document, or require further consultation with your Mirrored Disk/iX implementation, contact Roy Bittinger at 714-758-5319 or contact your local Hewlett-Packard Sales Office and ask for a Professional Services Technical Consultant.

2.0 REQUIREMENTS

2.1 PRODUCT CONFIGURATION REQUIREMENTS

Mirrored Disk/iX is designed to work with user volumes on Hewlett-Packard HP 3000 Series 900 systems. There are several configuration requirements:

- MPE software release A.30.00 or greater.
- Disk drives that use HP-FL cards or SCSI cards.
- Mirrored partners must be the same model of fiber-link drive or NIO SCSI drive.
- Mirrored partners must be connected to different HP-FL cards or SCSI controller cards.
- Mirrored Disk/iX only supports user volumes. Mirroring of the system volume set is not supported.
- Mirrored Disk/iX requires the use of user volumes for all data being mirrored.

2.2 USER CAPABILITIES

To initialize mirrored volumes system managers must have the **Create Volumes (CV)** capability. This capability is required in order to use VOLUTIL.

When an application writes to a file, Mirrored Disk/iX causes the file system to write that transaction to both partners simultaneously. When an application reads from a file, the read uses an algorithm that selects the drive that can respond fastest to the request.

2.3 PLANNING AND PREPARATION

The implementation of Mirrored Disk/iX requires some upfront planning and preparation in several areas. These include:

- High Availability for the MPEXL_SYSTEM_VOLUME_SET disks (system disks)
- Logical data mapping to user volumes
- Physical hardware, disk, and disk cabinet layout and configuration
- Implementation Planning
- Implementation Checklist
- Business Requirements of the Data

2.3.1 Planning High Availability for the MPEXL_SYSTEM_VOLUME_SET

Mirrored Disk/iX *does not* allow mirroring of any disks in the MPEXL_SYSTEM_VOLUME SET. This means that even if you protect all your other data with Mirrored Disk/iX, should one of the system disks fail (e.g., ldev 1 or 2) your system will go down. If this failure is caused by a bad disk which must be replaced, an INSTALL is required. All data from the system volume set must be re-installed.

There are basically three choices that you can make with regard to protecting your system volume set from the "disaster" described above:

1. **Use regular drives for the system volume set.**
"Live with the consequences if one of them fails". On a large MPE system (995 4-way), an *install* of MPE could take anywhere from 1.5 to 2 hours. This time doesn't include the time it takes for an HP CE to come on-site, remove the failed disk, install the new disk, run diagnostics, and configure the new disk into the system. The next step includes *restore* time for spool files and other files. In total, your system could be down from 6 to 8 hours should a system disk fail. If your company can afford to have the system down for that period of time, at any unpredictable moment in time, then you don't need to worry about protecting your system disks.

2. **Purchase at least two high availability (HA) disk arrays to use as ldevs 1 and 2.**
A disk array is actually a mechanism made up of multiple disks. High availability disk arrays are designed to increase data availability, increase total storage capacity, and provide flexibility. High availability disk arrays currently come in 4 giga-byte (GB) and 8GB flavors on the HP3000. The 4GB disk array is actually comprised of three 2GB disks, while the 8GB disk array is comprised of five 2GB disks. Data is spread across these disks. On the 4GB disk array, 2 disks are data disks and 1 disk is a parity disk. On the 8GB disk array, 4 disks are data disks and 1 disk is a parity disk. If any single disk within the disk array fails, the remaining disks are used to recreate the data in real time such that the data is still accessible. There is no interruption in the system processing. Failed disks can be replaced while the system is running and will automatically synchronize themselves.

   Recommendation:
   One potential disadvantage of currently available disk arrays is that each array acts as a single disk as far as disk I/Os are considered. There is only one path (i.e. one controller) and one disk head access at a time to service disk I/O requests. Using one disk array as the system disk has resulted in system performance bottleneck problems for some customer environments, so it is recommended that at least two be purchased. Of course, it always depends on your individual environment, however, you may not have time to test it out to see whether one array is enough. Typically, two 4GB disk arrays are sufficient to meet system disk requirements. The MPE/iX 5.5 release will introduce support for the F/W HA Disk Array which holds four half-height disks totaling 16GB, when using the new 4GB half-height disks (4x4). When the low profile disks are released, we will see these disk arrays support 32GB (8x4).

3. **Purchase additional regular disks, configure them as hot stand-by disks to the system disks.**
   This option has several very important limitations. For this reason, we do not recommend this option, however, we will explain what is involved should this be your only alternative.

   - If a system disk fails, the system will go down.
   - You have to shut the system down and do an INSTALL to the stand-by system disks to load them.
   - Any time system directory (account and group) changes are made, configuration changes, or MPE configuration changes are made, a System Load Tape (SLT) must be made and an INSTALL must be done to the stand-by disks as soon as possible to keep them in sync with the real system disks.
   - If a primary system disk fails, you switch to the stand-by disks, and the configuration of the system has changed but was not updated on the stand-by disks, the configuration changes made earlier are lost.

- When the replacement disk is installed, you must do an INSTALL to the primary system disks.

  This option is *not* recommended because of the operational complexities involved as well as the time required to maintain the configuration and recover from a disk failure.

2.3.2 Planning Logical Data Mapping to User Volumes

If your MPE/iX system is currently using only the MPEXL_SYSTEM_VOLUME_SET volume set, the next step in planning for Mirrored Disk/iX implementation is to determine how you will map your Accounts and Groups to user volumes. If you are already on user volumes, you may not need to perform this step.

Because Mirrored Disk/iX requires the use of user volumes, several key definitions on what user volumes are and how they work may be helpful.

2.3.2.1 Definitions:

The basic definitions required when discussing user volumes:

Volume A volume is a disk.

Volume Set A volume set is a set of disks that have been logically grouped together

Volumes A volume is a disk pack. A disk pack can be removable or non-removable. However, in today's world of 1, 2 and 4GB SCSI disks, we don't see removable packs anymore.

Volume Sets A volume set is a group of volumes containing one master volume and optional member volumes. There are two types of volume sets available on the system:

1) System volume set (MPEXL_SYSTEM_VOLUME_SET) that is initialized when the system is installed
2) Nonsystem volume sets that you create with VOLUTIL. You can create nonsystem volume sets and add volumes to volume sets while the system is running.

2.3.2.2 Sizing:

Before configuring user volumes, you must make the following sizing decisions:

- How much disk space do you need in the MPEXL_SYSTEM_VOLUME_SET (system volume set)
- What Groups and Accounts are going to map to which user volume sets
- How much disk space do you need per user volume set (2.3.2.4)
- How much free space do you need to allow on the user volume sets

How much disk space do you need in the MPEXL_SYSTEM_VOLUME_SET?
To calculate the amount of space you need on the system volume set, you have to know the following:

1. How much space does MPE/iX take?
   Typically this will be:
   a) SYS account
   b) all of the HP@ accounts (including HPSPOOL)
   c) any networking product accounts
   d) the TELESUP account
   e) third party products which must be installed on the system volume set
      (this includes software that is automatically tied to all data such as SharePlex/iX or Omnidex)
   f) how much growth do you expect?

2. Use the *REPORT* command to report current usage (in sectors) and then convert from sectors to Megabytes by multiplying by 256 and dividing by 1,000,000.
   Example: If your total of the account file space used from the *REPORT* command is 10,250,000 sectors, calculate Megabytes as follows:
   (10,250,000 x 256)/1,000,000=2,624 Megabytes or 2.624 Gigabytes

3. How much transient space does your **ENTIRE** system need?
   With the single MPEXL_SYSTEM_VOLUME_SET environment where user volumes are not utilized, transient space could be spread across all disks (or as many as you wanted). With user volumes configured, transient space can only be allocated on the system volume disks. What you really need to be able to do is accurately predict the high water mark of transient space usage on your system during production. One way to do this is to run HP GlancePlus, or other performance measurement application, periodically in batch mode at peak periods of processing during the day and during batch production at night. Save the STDLISTs and review them looking for the highest transient space usage. Than add at least 20% as a safety margin.

4. How much free space do you need to have on the system volume set?
   This will depend on the accounts and applications you load on the system volume set. Don't forget to allow space for log files (NMLG#### and LOG####). Again, be sure to add at least 20% or more as a safety margin.

5. How much space does the spooler use?
   Again, you want a high-water mark so that you allow enough space to handle all the files being created in the HPSPOOL account. Using the *REPORT* command periodically in batch during peak on-line and batch production periods can give you this information.

6. How many disk drives should be used for the system volume set?
   Once you've calculated the disk space you require for file space, system-wide transient space, spooler space, and free space, the next step is to determine how many disks to use in the system volume set. At this time you need to consider disk I/O performance on the system disks. You may find that disk space requirements only result in the need for one 2GB disk. However, experience has shown that putting all of MPE/iX on a single spindle (along with transient space and spool space) has a serious effect on overall system performance because of the high number of physical disk I/Os performed on the system volume set. Consequently, it is recommended that you still provide at least two disks (three are better) for the system volume set.

Which Groups and Accounts Map to Which User Volume Sets
Determining which Groups and Accounts to put on user volume sets and how many user volume sets may be necessary is difficult. For example, two approaches that can be considered include:

1. Put all data that does not belong on the system volume set on one single user volume set. This makes disk space management and system administration easier but eliminates one of the advantages of user volumes. An advantage of user volumes is that should both partners of a mirrored pair fail, the user only has to restore the data from one user volume set rather than restoring the whole system. The larger the user volume set the more time it will take to restore it in the event of a catastrophe.

2. Logically divide the data into volume sets by *ACCOUNT* (user volumes actually assign files to volume sets by *GROUP* but it can be very confusing having two *GROUP*s of the same *ACCOUNT* on different volume sets). The following is an adaption of the volume set arrangement being used by a Southern California shoe manufacturer and distributor:

Volume Set Name	Contents
ORDERS_VOLUME_SET	All data and programs involved in the Order Entry applications
DISTRIB_VOLUME_SET	All data and programs involved in the Distribution applications
FIN_VOLUME_SET	All data and programs involved in the Financial applications
UTIL_VOLUME_SET	All third party utilities (MPEX, Adager, etc.)

The customer chose this volume set structure because it made logical sense along account boundaries and kept the volume set sizes reasonable (in their case, none are larger than 12 Gigabytes). Also, each application has its own schedule so backups are done at different times.

2.3.2.3 Calculating Disk Requirements for Each User Volume Set

Once you've chosen a logical division of accounts and assigned them to volume sets, you can calculate the disk space requirements for each user volume set.

1. Include 100 Mbytes on the first disk of each volume set. This space is used by the MPE/iX Transaction Manager and other MPE/iX data structures - including the Volume Table and Transaction Log Files.

2. Run *REPORT* on each account that is to be located on the user volume set. This allows for the calculation of disk space requirements for the user volume set.

3. The last step is to calculate the free space needed on each user volume set.

How much free space do you need to allow for on the user volume sets?
Calculating the free space needed for each user volume set can be a little tricky. Remember that you need to allow room for temporary files, sort files, and file growth on each user volume set. Some customers have undersized the free space in a user volume set because they forgot to include the sort files that their applications created during production or the file space needed by Adager or DBGeneral to expand IMAGE data sets. Remember that if a program creates a file (permanent or temporary) in a group that is assigned to a user volume set, there has to be room on that user volume set to build that file or the application will abort. If you are not sure, it's better to overestimate. In most cases, you will use the disk space eventually.

2.3.3 Physical Hardware, Disk, and Disk Cabinet Layout and Configuration

Now that you know what your system volume set will look like, how many user volume sets you will have, and how much disk space you need in user volume set, you need to determine how many disks will be needed to implement Mirrored Disk/iX, how they should be installed in cabinets and how to connect them to the system. The following is a list of steps you will need to complete:

1. Determine how many and which types of disks will be in each volume set
2. Determine how many and which types of disks are needed to mirror each volume set
3. Determine how many interface cards are needed
4. Determine how many slots are needed
5. Determine the best way of arranging the disks in cabinets
6. Determine the number of cabinets and cables needed

1. **Determine how many and which types of disks will be in each volume set.**
Each user volume set must be made up of one or more dedicated disk drives. A disk drive can only be a member of one volume set. For each user volume set identified above, determine the number of disks required based on the data space and free space requirement estimates made above. System and application performance need to be considered at this point. If your system is currently experiencing data contention problems for data that is to be located on a user volume set, it is better to increase the number of disk drives by either choosing smaller capacity drives or merely increasing the number of drives. For most applications that are based on IMAGE data bases, you will want to spread associated master and detail data set files on separate disks. This may increase the number of drives you will need for the volume set.

   Create a diagram for each volume set of the disks to be assigned to that volume set and the major data files which need to be placed on each disk for performance purposes.

2. **Determine how many and which types of disks are needed to mirror each volume set.**
Once the number and type of disks have been determined for the user volume set primary disks, determining the number and type of disks for the mirrored partners is simple. Mirrored partners must be exactly the same disk model type and each primary disk must have one mirrored partner disk.
Example: If the user volume set will need five 2GB disk drives for the primary disks, then specify five 2GB disk drives for the mirrored partner disks.

3. **Determine how many interface cards are needed.**
Mirrored Disk/iX requires that the mirrored partner disks be on separate interface cards than the primary disks. This means that each user volume set will require an even number of interface cards. The number of cards required is influenced by the following guidelines:

   A. SCSI II cards support a maximum of 7 disks per chain but for performance reasons it is better not to configure more than 5 per chain.
   B. Fast and Wide SCSI cards support a maximum of 15 disks per chain. Again, for performance reasons it is better not to configure more than 10 disks per chain.
   C. You shouldn't share a SCSI or FL card between the MPEXL_SYSTEM_VOLUME_SET and a mirrored user volume set.
   D. Multiple user volume sets can be on the same SCSI chain. However this does have a consequence if a disk on the chain fails and has to be replaced. To replace a SCSI disk, all disks on the chain must be powered down. This means that any volume sets represented by disks on the chain will have to operate in a non-mirrored state while the disk is being replaced. After the disk is replaced, all disks on the chain will automatically go into repair state (Mirrored Disk/iX copies data to each of the disks on the chain from each disk's mirrored partner). This will not affect application performance or result in data loss, it will just keep the disks on the chain fairly busy until repair completes. If the plan is to shut the system down completely any time a disk needs to be replaced, sharing a SCSI card between multiple user volume sets has no drawbacks.

E. Choosing between SCSI II and Fast and Wide SCSI may or may not be a difficult choice. SCSI II allows the system to transfer up to 5 Megabytes per second of data to the system's internal I/O bus. Fast and Wide allows the transfer of up to 20 Megabytes per second. SCSI II requires a half-height card. Fast and Wide requires a full-height card. You will need to factor in the existing disks and I/O cards you already have as well as the number of half-height slots versus full-height slots available on the system. Just remember that mirrored partners must be identical discs on the exact same type of I/O card. With SCSI disks, the partners must have the same unit number also (i.e. if the path to the disk is 52.6.1, then 1 is the unit number).

4. **Determine how many slots are needed.**
The number of I/O slots needed is determined by the number and type of I/O cards required. You may have to adjust the type of I/O cards and number of chains based on the number of slots you have on your system. In many cases, implementing Mirrored Disk/iX requires upgrading to a larger chasis because of a shortage of I/O slots. It is often advisable to try to size your system so that you have at least 2 I/O slots open to allow for future growth.

5. **Determine the best way of arranging the disks in cabinets.**
Because mirrored disk configurations can get complicated and usually include a large number of disk drives, it is best to arrange the disks in cabinets so that it is easy to determine which physical disks are mirrored partners and where the disks of each volume set reside. There are many ways to arrange the cabinets. This becomes especially critical if a disk needs to be replaced. It is important to ensure that the correct disk is replaced and that it's mirrored partner is easy to identify. If a replaced disk is identified with the wrong partner at replacement time (with the REPLACEMIRRVOL command) data can be lost!

To help the Customer Engineers (CE) install the mirrored disks and to help operations during a crisis, you should create diagrams for each system showing all cabinets, the disks in each cabinet, and which disks are mirrored pairs. This diagram is also beneficial should you decide to change your configuration or add-on.

The following is an example of the disk and cabinet layout being used by an insurance company in Southern California:

This company has an application that shares one very large data base. The company already had a number of 2 GB SCSI II disks and cards and did not want to invest in disk arrays for ldevs 1, 2, and 3.

Ldevs 1, 2, and 3 are 2 GB system disks.
Ldevs 22 through 29 are Compressed DAT Tape Drives.
Ldevs 4 and 5 are disks reserved for Auto Restart Dump-to-Disk.
Ldevs 30 though 50 are primary disks for PROD_VOLUME_SET (which is where all the customer's application data resides).
Ldevs 51 through 53 are primary disks for UTIL_VOLUME_SET (where all third party utilities reside).
Ldevs 60 through 80 are the mirrored partners of ldevs 30 through 50. Each partner is in the same physical location but in separate cabinets. This makes it simple to identify not only which disk needs to be replaced but which physical disk is that disk's mirrored partner.
Ldevs 81 through 83 are the mirrored partners of ldevs 51 through 53.
Ldevs 801 through 803 are system volume hot spare disks.

6. **Determine the number of cabinets and cables needed.**
The number of cabinets and SCSI cables needed depends completely on how you decide to lay out the cabinets. It is recommended that an HP Customer Engineer assist in determining the number of cabinets and cables you will need because of SCSI cabling restrictions, power, and cabinet racking restrictions.

2.3.4 Implementation Planning

Planning the actual implementation of Mirrored Disk/iX must be done carefully and thoroughly. Remember that when implementing Mirrored Disk/iX you are essentially wiping all data off of your system (including the directory) and reloading it. The following is an outline of the steps required during planning:

1. Determine which files need to be placed on which disks in each user volume set.
2. Use BULDACCT.PUB.SYS to create jobs that will build the new directory for each user volume set.
3. Create custom store jobs for each volume set
4. Create custom restore jobs for each volume set
5. Make sure you have GOOD tape media (DAT, reel-to-reel, etc.).
6. Create an alternate config group with SYSGEN representing the new I/O config of the system after all disks are installed.

1. Determine which files need to be placed on which disks in the user volume set.
This is only needed to help avoid system performance problems due to disk contention after the full restore is done. If you already have restored certain data sets to certain disks for performance reasons on your existing system, you will need to determine where these data sets need to go in the user volume set. Make a list of the files that need to go on each disk in the volume set.

   This one step can be *especially crucial* to system performance. In several instances, merely restoring all data in a list of accounts to a user volume set has resulted in a few of the disks in the volume set being heavily used and the rest to be lightly accessed. This results in long wait-for-disk queues which slows down application performance. The best approach is to try to spread your busiest files across separate disks so that disk I/Os are spread more evenly throughout the user volume set. You need to determine which files need to be on which disks **BEFORE** you create your store tapes.

2. Use BULDACCT.PUB.SYS to create jobs that will build the new directories for each user volume set.
BULDACCT is a utility supplied by HP as part of MPE which can be used to create jobs that will create all the entries needed in both the system directory (on ldev 1) and the user volume directory (on the first disk, or master disk, of each user volume set). When BULDACCT is run, it creates two jobs, BULDJOB1 and BULDJOB2. BULDJOB1 contains all the NEWACCT, ALTACCT, NEWGROUP, ALTGROUP, NEWUSER, and ALTUSER commands needed to create the directories. BULDJOB2 sets up the UDCs for each user. BULDACCT has many options. BULDACCT will need to be run for each user volume set as well as for the system volume set. After each run, you need to rename the BULDJOB1 and BULDJOB2 files so that they don't get overwritten by the next run of BULDACCT.

Example: If you are moving from the MPEXL_SYSTEM_VOLUME_SET environment, use the following BULDACCT command to create BULDJOB1 and BULDJOB2 for the user volume set "PROD_VOLUME_SET" containing the following user accounts: DATA, DOC, LIB, PROG

```
:RUN BULDACCT.PUB.SYS;
INFO="DATA,DOC,LIB,PROG%VSACCT=PROD_VOLUME_SET"

RENAME BULDJOB1,BLDPROD1.BULDJOB
RENAME BULDJOB2,BLDPROD2.BULDJOB
```

To run *BULDACCT* for the system volume set containing the HP@, HPSPOOL, SYS, and TELESUP accounts use the following command:

```
:RUN BULDACCT.PUB.SYS;INFO="HP@,SYS,TELESUP"
RENAME BULDJOB1,BLDSYS1.BULDJOB
RENAME BULDJOB2,BLDSYS2.BULDJOB
```

3. Create custom store jobs for each volume set.
The most time consuming part of implementing Mirrored Disk/iX is storing and restoring data from and to the disks. Because the entire disk layout and contents is being replaced, every byte of data (data, programs, utilities, etc.) will have to be restored to the system. To shorten this task, you should create individual stores for each type of restore you will need to do after configuring the volume sets. You should also try to make use of all the tape drives you have on the system so that you can do parallel restores. In addition, you may want to restore certain files (or IMAGE data sets) to certain ldevs. Create separate jobs to store these files on separate tapes, 1 set of tapes for each individual disk drive that you are going to specifically restore to. This will allow you to restore data quickly and accurately.

Example: You have decided to configure four volume sets on your system:

MPEXL_SYSTEM_VOLUME_SET	- containing SYS, HP@, and TELESUP accounts
PROD_VOLUME_SET	- containing the programs and data for a manufacturing application
FIN_VOLUME_SET	- containing the programs and data for the financial applications
UTIL_VOLUME_SET	- containing the third party utilities

(Your system has four DDS tape drives)

Create the following store jobs and put them in the **STRJOB.SYS** directory:

1. BULDSTR to store all BLD@.MIRRJOB.SYS jobs: BLDSYS1, BLDSYS2, BDPROD1, BLDPROD2, BLDFIN1, BLDFIN2, BLDUTIL1, BLDUTIL2.
2. MPESTR1 to store @.PUB.SYS, @.DIAG.SYS-LOG@.PUB.SYS-NMLLG@.PUB.SYS (these files will all be restored to ldev 1)
3. MPESTR2 to store @.@.SYS,@.@.TELESUP,@.@.HP@-@.PUB.SYS-@.DIAG.SYS-LOG@.PUB.SYS-NMLLG@.PUB.SYS (these files will be spread across all mpe system disks)
4. PRDSTR31 to store certain manufacturing files that will be restored to ldev 31
5. PRDSTR32 to store certain manufacturing files that will be restored to ldev 32
6. PRDSTR33 to store certain manufacturing files that will be restored to ldev 33
7. PRDSTR34 to store certain manufacturing files that will be restored to ldev 34

8. PRDSTRA to store 50% of the data from the manufacturing account data not already accounted for by the specific stores.
9. PRDSTRB to store the other 50% of the manufacturing account data not already accounted for by the specific stores.
10. FINSTR35 to store certain financial files that will be restored to ldev 35
11. FINSTR36 to store certain financial files that will be restored to ldev 36
12. FINSTR37 to store certain financial files that will be restored to ldev 37
13. FINSTR38 to store certain financial files that will be restored to ldev 38
14. FINSTRA to store the first 50% of the financial data not already stored on the tapes above.
15. FINSTRB to store the other 505% of the financial data not already stored on the tapes above.
16. UTILSTR1 to store 50% of the utility account data.
17. UTILSTR2 to store the other 50% of the utility account data.
18. ALLSTR1 to store the first 25% of all accounts (full backup).
19. ALLSTR2 to store the second 25% of all accounts.
20. ALLSTR3 to store the third 25% of all accounts.
21. ALLSTR4 to store the fourth 25% of all accounts.

These store jobs assume that the PROD_VOLUME_SET and FIN_VOLUME_SET are going to be fairly large (more than 4 GB each) and that the UTIL_VOLUME_SET is fairly small (around 2 GB or less). The idea is to set up a store that creates a tape that can be restored in parallel with other tapes and allow the entire system to be restored as quickly as possible instead of one massive restore with all files on one tape volume. In addition, this method will allow you to reduce the initial store time.

The ALLSTR jobs would do a complete system backup of all data. This is a good insurance policy in case one of the specific store tapes is bad, gets lost, gets destroyed, etc.

Be sure to validate each tape (VSTORE seems to work fairly well). This is critical. Remember that once mirrored disks have been configured, all user data is erased from the disks and has to be restored from tapes. If you have a bad tape, you will lose data.

4. Create custom restore jobs for each volume set.
The next step is to create custom restore jobs for each store job. You want to do this ahead of time so that during the heat of implementation, a set of files is not left out. The restore jobs should mirror the store jobs (there should be one restore job for each store job).

Example: In step 3 above, you created 14 store jobs. Create a restore job for each store job and put it in the **RSTJOB.SYS** directory:

1. BLDRST to restore all BLD@.MIRRJOB.SYS jobs: BLDSYS1, BLDSYS2, BLDPROD1, BLDPROD2, BLDFIN1, BLDFIN2, BLDUTIL1, BLDUTIL2
2. MPERST1 to restore
@.PUB.SYS, @.DIAG.SYS-LOG@.PUB.SYS-NMLLG@.PUB.SYS
(these files will all be restored to ldev 1)
3. MPERST2 to restore
@.@.SYS,@.@.TELESUP,@.@.HP@-@.PUB.SYS-@.DIAG.SYS-LOG@.PUB.SYS-NMLLG@.PUB.SYS;KEEP;OLDDATE;CREATE
(these files will be spread across all mpe system disks)
4. PRDRST31 to restore certain manufacturing files to ldev 31
5. PRDRST32 to restore certain manufacturing files to ldev 32
6. PRDRST33 to restore certain manufacturing files to ldev 33
7. PRDRST34 to restore certain manufacturing files to ldev 34
8. PRDRSTA to restore 50% of the data from the manufacturing account data not already accounted for by the specific stores.
9. PRDRSTB to restore the other 50% of the manufacturing account data not already accounted for by the specific stores.

10. FINRST35 to restore certain financial files to ldev 35
11. FINRST36 to restore certain financial files to ldev 36
12. FINRST37 to restore certain financial files to ldev 37
13. FINRST38 to restore certain financial files to ldev 38
14. FINRSTA to restore the first 50% of the financial data not already restored from the tapes above.
15. FINRSTB to restore the other 50% of the financial data not already restored from the tapes above.
16. UTILRST1 to restore 50% ot the utility account data
17. UTILRST2 to restore the other 50% of the utility account data
18. ALLRST1 to restore the first 25% of all accounts
19. ALLRST2 to restore the second 25% of all accounts
20. ALLRST3 to restore the third 25% of all accounts
21. ALLRST4 to restore the fourth 25% of all accounts

Now create a tape of all the RST@.RSTJOB.SYS files and validate it.

5. Make sure you have GOOD tape media (DAT, reel-to-reel, etc.)
   Before actual implementation, make sure you have a more than adequate supply of good tapes. DAT and reel-to-reel tapes cannot always be depended on to be good. If you don't have known, good tapes, either purchase additional tapes, or test each tape before it will be used.
   You cannot afford to run into tape media problems during the implementation.

6. Create an alternate config group with SYSGEN representing the new I/O config of the system after all disks are installed.
   While this step is not required, it will save a lot of time at implementation time. MPE allows you to create an SLT tape with a different configuration on it without affecting the current system. In fact, MPE allows you to have multiple I/O configurations on a system at one time (but only one is active at a time). The default I/O configuration is kept in the group CONFIG.SYS. In this step, have your HP CE go into SYSGEN and modify the I/O configuration so that it is the configuration that you want after all the disks are installed and connected. This may take 30 minutes to an hour. Save this configuration into a different config group (such as MIRRCFG.SYS). At implementation time, after all disks are installed and connected, you can tell MPE to use this alternate config group as part of the startup. This will allow you to save the 30 minutes to an hour of CE time during your implementation window and will also avoid potential data entry mistakes.

   The following is the syntax the CE will use to save the config group MIRRCFG:
   Make sure you store this group @ .MIRRCFG.SYS on to tape before you begin implementation.

2.3.5 Implementation

The implementation of Mirrored Disk/iX is fairly complicated and lengthy. The following is an outline of the steps required actually implement Mirrored Disk/iX:

```
sysgen> io
io> ... CE makes all the additions, modifications,
        and deletions
io> hold
io> exit
sysgen> keep MIRRCFG
sysgen> exit
```

1. Verify that all disks, cables, cabinets, and software have arrived
2. Install Mirrored Disk/iX software
3. Install any patches required by Mirrored Disk/iX

4. Perform all data stores, backups and tape validation
5. Shutdown the system
6. HP CE installs disks, cabinets, and cables
7. Label all disks and cabinets with ldev#s
8. INSTALL from the SLT tape
9. RESTORE the new configuration group
10. Point to the new configuration group in SYSGEN
11. Restart the system
12. Show the status of the drives and verify
13. Restore the BULDACCT jobs to ldev1
14. Create the MPEXL_SYSTEM_VOLUME_SET accounting structures
15. Restore the MPESTR1 files to Ldev 1
16. Add the other system disks to MPEXL_SYSTEM_VOLUME_SET
17. Restore the MPESTR2 files to the MPEXL_SYSTEM_VOLUME_SET
18. Create the PROD_VOLUME_SET mirrored configuration
19. Create the FIN_VOLUME_SET mirrored configuration
20. Create the UTIL_VOLUME_SET mirrored configuration
21. VSCLOSE and VSOPEN each volume set
22. Create the PROD_VOLUME_SET accounting structures
23. Create the FIN_VOLUME_SET accounting structures
24. Create the UTIL_VOLUME_SET accounting structures
25. Restore the PROD_VOLUME_SET data
26. Restore the FIN_VOLUME_SET data
27. Restore the UTIL_VOLUME_SET data
28. Enable system-wide UDC's
29. Enable PROD_VOLUME_SET UDC's
30. Enable FIN_VOLUME_SET UDC's
31. Enable UTIL_VOLUME_SET UDC's
32. Test network connections
33. Test applications and data access
34. Start up production

Detailed explanation of the above steps:

1. Verify that all disks, cables, cabinets, and software have arrived.

2. Install Mirrored Disk/iX software
   Use the instructions that come with the Mirrored Disk/iX software to install it. To check that Mirrored Disk/iX is installed, run the VOLUTIL program. It should come up with the banner 'MIRVUTIL'. This indicates that Mirrored Disk/iX software is installed.

3. Install any patches required by Mirrored Disk/iX
   Call the Response Center or check with your Software Support Engineer to get the latest patches (if any) for Mirrored Disk/iX. Chances are that the latest Express Patch has any patches you might need. However, even if you already applied it, you will have to apply the patches again because you didn't have Mirrored Disk/iX installed.

4. Perform all data stores and backups
   This step is absolutely critical. You must make sure that all files are backed up and that all backup jobs complete successfully. You also should verify that the tapes are readable. The VSTORE utility in PUB.SYS works well at verifying the readability of tapes. In the example above, you would want to make sure that you have successfully created and verified the sets of tapes described in the Mirrored Disk/iX Implementation Store Tape Checklist.

5. Shutdown the system
   BEFORE shutting down the system, make sure all your tapes are clearly labeled and all have been verified!. Remember, once the CEs have moved disks around and installed new disks, your data is gone!!!

6. HP CEs install disks, cabinets, and cables.
   This could take a few hours depending on how many disks and cabinets you have.

7. Label all disks and cabinets with ldev#s
   It is recommended that each component in each cabinet be labelled with the ldev# of the disk/tape drive. This way, it will be clear which drive is down and easy to find it's partner.

8. INSTALL from the SLT tape
   Mount CSLT tape and on the console type:

CTRL-B
CM>**RS**

```
ENTER Y TO CONFIRM YOUR INTENTION TO RESTART THE SYSTEM (Y/N)? Y
```

Watch console messages and look for:

```
Autoboot from primary path enabled
To override, press any key within 10
```

YOU MUST PRESS ANY KEY ON KEYBOARD NOW!

```
BOOT FROM PRIMARY BOOT PATH (Y or N)? N
BOOT FROM ALTERNATE BOOT PATH (Y OR N)? Y
INTERACT WITH IPL (Y or N)? Y
```

Wait for the 'Booted' message and then the ISL prompt; type install, and start norecovery nosysstart logon=manager.sys

```
ISL>INSTALL
ISL>START NORECOVERY NOSYSSTART LOGON=MANAGER.SYS
```

Wait for the system to come up...

9. RESTORE the new configuration group

```
FILE T;DEV=TAPE
RESTORE *T;@.MIRRCFG.SYS;OLDDATE;SHOW;CREATE
```

10. Point to the new configuration group in SYSGEN

```
:SYSGEN
sysgen>ba MIRRCFG.SYS
sysgen>keep CONFIG.SYS
sysgen>exit
```

11. Restart the system

    CTL-A SHUTDOWN
    CTL-B RS

```
ENTER Y TO CONFIRM YOUR INTENTION TO RESET THE SYSTEM (Y/N)? Y
```

Watch console messages and look for:

```
Autoboot from primary path enabled
To override, press any key within 10
```

Allow system to boot from Primary Path

```
INTERACT WITH IPL (Y or N)? Y
```

Wait for the 'Booted' message and then the ISL prompt; type start norecovery.

```
ISL>START NORECOVERY
```

Wait for system to come up...

12. Show the status of the drives and verify

```
:DSTAT ALL
```

Any drives which do not come up in SCRATCH, UNKNOWN, or ERROR MOUNT state must be scratched using the SCRATCHVOL command:

```
 :VOLUTIL
volutil:SCRATCHVOL 2 (scratch ldev 2)
```

13. Restore the BULDACCT jobs to ldev1

```
:STREAM BLDRST
```

This job restores the following files:
BLDSYS1, BLDSYS2
BLDPROD1, BLDPROD2
BLDFIN1, BLDFIN2
BLDUTIL1, BLDUTIL2

14. Create the MPEXL_SYSTEM_VOLUME_SET accounting structures

```
: STREAM BLDSYS1
```

15. Restore the MPESTR1 files to Ldev 1
First restore all the RSTJOB restore jobs:

```
:FILE T;DEV=TAPE
:RESTORE *T;@.RSTJOB.SYS;OLDDATE;CREATE;KEEP
```

Now stream the job to restore MPE files to ldev 1

```
:STREAM MPERST1
```

16. Add the other system disks to MPEXL_SYSTEM_VOLUME_SET.
Run DSTAT, scratch drives if necessary and configure system drives

```
:DSTAT ALL
```

Any drives which do not come up in SCRATCH, UNKNOWN, or ERROR MOUNT state must be scratched using the SCRATCHVOL command:

```
:VOLUTIL
volutil: SCRATCHVOL 2 (scratch ldev 2)
```

Configure system drives:

```
volutil: NEWVOL MPEXL_SYSTEM_VOLUME_SET:MEMBER22 100 100
volutil: NEWVOL MPEXL_SYSTEM_VOLUME_SET:MEMBER33 100 100
```

Now verify that the MPEXL_SYSTEM_VOLUME_SET is configured correctly:

```
volutil: DSTAT ALL
DSTAT should show:

Ldev 1 as MASTER
Ldev 2 as MEMBER2
Ldev 3 as MEMBER3
```

All other Ldevs should show as SCRATCH or UNKNOWN

17. Restore the MPESTR2 files to the MPEXL_SYSTEM_VOLUME_SET

```
:STREAM MPESTR2.RSTJOB.SYS
```

18. Create the PROD_VOLUME_SET mirrored configuration

```
          :VOLUTIL
volutil:: NEWMIRRSET PROD_VOLUME_SET MEMBER1 (31,61)
volutil:: VSCLOSE PROD_VOLUME_SET
volutil:: VSOPEN PROD_VOLUME_SET
volutil:: NEWMIRRVOL PROD_VOLUME_SET: MEMBER2 (32,62)
volutil:: NEWMIRRVOL PROD_VOLUME_SET: MEMBER3 (33,63)
volutil:: NEWMIRRVOL PROD_VOLUME_SET: MEMBER4 (34,64)
volutil:: SHOWSET PROD_VOLUME_SET  MIRROR
```

19. Create the FIN_VOLUME_SET mirrored configuration

```
          :VOLUTIL
volutil:: NEWMIRRSET FIN_VOLUME_SET MEMBER1 (35,65)
volutil:: VSCLOSE FIN_VOLUME_SET
volutil:: VSOPEN FIN_VOLUME_SET
volutil:: NEWMIRRVOL FIN_VOLUME_SET: MEMBER2 (36,66)
volutil:: NEWMIRRVOL FIN_VOLUME_SET: MEMBER3 (37,67)
volutil:: NEWMIRRVOL FIN_VOLUME_SET: MEMBER4 (38,68)
volutil:: SHOWSET FIN_VOLUME_SET  MIRROR
```

20. Create the UTIL_VOLUME_SET mirrored configuration

```
          :VOLUTIL
volutil:: NEWMIRRSET UTIL_VOLUME_SET MEMBER1 (39,69)
volutil:: NEWMIRRVOL UTIL_VOLUME_SET: MEMBER2 (40,70)
volutil:: VSCLOSE UTIL_VOLUME_SET
volutil:: VSOPEN UTIL_VOLUME_SET
volutil:: SHOWSET UTIL_VOLUME_SET  MIRROR
```

22. Create the PROD_VOLUME_SET accounting structures

```
:STREAM 10
:STREAM BLDPROD1
```

23. Create the FIN_VOLUME_SET accounting structures

```
:STREAM BLDFIN1
```

24. Create the UTIL_VOLUME_SET accounting structures

```
:STREAM BLDUTIL1
```

25. Restore the PROD_VOLUME_SET data

Use as many tape drives as possible to perform parallel stores. This way the production system will be ready sooner.

```
:STREAM PRDRST31 to restore certain manufacturing files to ldev 31
:STREAM PRDRST32 to restore certain manufacturing files to ldev 32
:STREAM PRDRST33 to restore certain manufacturing files to ldev 33
:STREAM PRDRST34 to restore certain manufacturing files to ldev 34
:STREAM PRDRSTA to restore 50% of the data from the manufacturing
   account data not already accounted for by the specific stores
:STREAM PRDRSTB to restore the other 50% of the manufacturing
   account data  not already accounted for by the specific stores
```

26. Restore the FIN_VOLUME_SET data

```
:STREAM FINRST35 to restore certain financial files to ldev 35
:STREAM FINRST36 to restore certain financial files to ldev 36
:STREAM FINRST37 to restore certain financial files to ldev 37
:STREAM FINRST38 to restore certain financial files to ldev 38
:STREAM FINRSTA to restore the first 50% of the financial data
     not already restored from the tapes above.
:STREAM FINRSTB to restore the other 50% of the financial data
     not already restored from the tapes above.
```

27. Restore the UTIL_VOLUME_SET data

```
:STREAM UTILRST1 to restore 50% to the utility account data
:STREAM UTILRST2 to restore the other 50% to the utility
   account data
```

28. Enable system-wide UDC

```
:STREAM BLDSYS2.BULDJOB.SYS
```

29. Enable PROD_VOLUME_SET UDC's

```
:STREAM BLDPROD2.BULDJOB.SYS
```

30. Enable FIN_VOLUME_SET UDC's

```
:STREAM BLDFIN2.BULDJOB.SYS
```

31. Enable UTIL_VOLUME_SET UDC's

```
:STREAM BLDUTIL2.BULDJOB.SYS
```

32. Test network connections

33. Test applications and data access

34. Start up production

2.3.5.1 Operational Procedure Changes

Many commands in MPE that are embedded in job streams or command files may need to be modified because of the implementation of user volume sets. For job streams or command files that execute on accounts or groups that are assigned to the MPEXL_SYSTEM_VOLUME_SET, no changes are needed.

The following commands may need to be changed:

- DISCFREE
- NEWACCT, ALTACCT, PURGEACCT, LISTACCT
- NEWGROUP, ALTGROUP, PURGEGROUP, LISTGROUP
- STORE
- RESTORE
- REPORT

The most common commands that need to be modified are those which create, alter, purge, or list accounts and groups. This is because with user volumes we now have multiple directories to maintain. The system directory still exists on ldev1. But in addition on the first volume of each volume set, there is a directory that represents the groups and files on that volume set. Consequently, adding, altering, or purging any group or account requires two steps, one to take care of the entry in the system directory on ldev 1 and one to take care of the user volume directory.

Example: To add the account MANU to the PROD_VOLUME_SET, the following commands must be entered:

```
:NEWACCT MANU,MGR;CAP=IA,BA,NS,SF,ND....
:NEWACCT MANU;ONVS=PROD_VOLUME_SET
```

To add a new group to a user volume set, two NEWGROUP commands must be entered:

```
:NEWGROUP DATA;CAP=.....;HOMEVS=PROD_VOLUME_SET
:NEWGROUP DATA;ONVS=PROD_VOLUME_SET
```

It is **CRITICAL** that all NEW/ALT/PURGE ACCT/GROUP commands conform to this two step process. If you are not careful, it is easy to alter a group so that the system directory points to the MPEXL_SYSTEM_VOLUME_SET instead of the user volume set. If this happens, it looks like all the data is gone. However, in reality, the data is still there, the directory is just pointing to the wrong volume set. It is recommended that use of the NEW/ALT/PURGE ACCT/GROUP commands be strictly limited so that production data is not "misplaced".

The utility volcheck.pub.sys is available from the Interex Users Group Contributed Library Tape. This utility will verify that all groups in an account are on the same volume set and report any groups accidentally put on the system volume set.

One other area for consideration is that any third-party software that is to be installed after the fact will probably install to the system volume set. You will need to modify the install jobs to create the accounts and groups in the user volume set desired.

Ongoing File Space Management:
User Volumes require that free space be managed by volume set. If a volume set fills up, the application will probably abort. The *DISCFREE* command will show free space by drive. However, it doesn't show space by volume set. The contributed library program VOLINFO.PUB.SYS will report free space by volume set. This utility will allow you to quickly see the disk usage and free space of each volume set.

Managing User Volumes:
The "Managing User Volume Sets" reference manual is a great source of operational tips for managing the user volumes you created for mirrored disks.

UV and CV Capability
Each account and user who will be creating groups on the system will need to have "UV" capability. Users who will be creating accounts will need "CV".

Integration of Disk Failure procedures into existing Operating Procedure

Before discussing the operational procedures required to support failure of mirrored disks, now is probably a good time to review how mirrored disk **recovery** works.

Disk repair is a mirrored disk operation that copies data from the good drive to the out-of-sync drive to bring a mirrored pair to a consistent state without interrupting applications accessing the volume set. After the repair operation is completed, normal mirroring resumes.

The system starts repairing a disk when one of the following occurs:

Operator starts repair	Operator issues the REPLACEMIRRVOL command to start repair.
System automatically starts repair	Upon volume mount, the system checks to make sure both partner disks contain the same information. If the disks were not VSCLOSE'd, the system selects, as the source, the drive that mounted first and repairs to the partner drive.

All mirrored pairs on the system cannot be repaired at once. There is a limit of six repair operations taking place at the same time. This is to limit the performance impact of repairing on the system. The repairs are staged, such that when one of the six repairs finishes, another begins.

Drives that are staged (awaiting repair) transition to the repair state in order of mounting. If the maximum number of repairs is taking place, subsequent pairs that mount and get staged are serviced in the order that they are mounted.

In the event of a system failure, Mirrored Disk/iX performs disk repairs automatically for any pair that is out-of-sync.

You can use the VOLUTIL SHOWSET command to display volumes involved in the repair process.

```
:VOLUTIL
  Mirvutil A.00.00, (C) Hewlett-Packard Co., 1990.
  All Rights Reserved.

  volutil:SHOWSET PROD_SET MIRROR

Volume Name   Vol Status   Mirr Status    Ldev   Mirr ldev
MEMBER1         MASTER       NORMAL        30       31
MEMBER1         MASTER       NORMAL        31       30
MEMBER2         MEMBER       REPAIR-DEST   32       33
MEMBER2         MEMBER       REPAIR-SRCE   33       32
```

The above screen shows that LDEV 32 (REPAIR-DEST) is being repaired by LDEV 33 (REPAIR-SRCE). The repair process takes about twenty minutes to complete. Programs and data residing on MEMBER2 are available while repairs are taking place.

Disk Failures

There are three types of disc failures that can occur for which operations has to be ready. Each requires a different recovery procedure:

- System disk (MPEXL_SYSTEM_VOLUME_SET) failure
- User volume disk failure during system startup, power on, or mounting (other than the master disk of the user volume set)
- User volume disk failure after system startup, power on, or mounting (other than the master disk of the user volume set)

The instructions that follow assume that you cannot 'hot swap' a failed disk drive (i.e. replace the failed drive while all other drives are up and running). This means that we assume that you will replace the disk under the safest conditions during planned downtime. That is, with MPE/iX down and power down on all drives.

❑ System disk failure

If a system disk fails at any time (and it is not a disk array), the disk will have to be replaced and an INSTALL will have to be done on the disks which are part of the MPEXL_SYSTEM_VOLUME_SET. The nice thing about this is that the data on the user volume disks is still intact. This implies that you should always have a current Custom System Load Tape (CSLT) ready along with backup tapes of the accounts that are located on the system disks and BULDACCT job files to rebuild ALL accounts, groups, and UDCs. The steps to be followed in this case are as follows:

1. Place a service call to HP and get the failed disk replaced and added into the MPEXL_SYSTEM_VOLUME_SET
2. Perform and INSTALL from the most recent CSLT tape
3. Restore the BULDACCT job files to the BULDJOB group of SYS
4. Rebuild the accounting structure for the system disks by streaming BLDSYS1
5. Stream the RSTSYS1 job to restore the MPE files to ldev 1
6. Stream the RSTSYS2 job to restore the rest of the system disk files

7. Rebuild the connections between the system directory and each user volume directory by streaming the BULDACCT jobs:
   BLDPROD1, BLDFIN1, BLDUTIL1
8. Reassign UDCs using the BULDACCT jobs
   BLDSYS2, BLDPROD2, BLDFIN2, BLDUTIL2

NOTE : When a disk in the system volume set fails, all current spoolfiles will be lost (unless they were processed and archived by a third party product such as SPOOLMATE). Operations must be ready to research and document those reports which were lost and prepare a recovery procedure for these.

❑ **User volume disk failure during system startup, power on, or mounting**
The system automatically mounts a mirrored volume set after volumes have been added to a mirrored volume set, upon power on of the disk drive, or the boot of the system. When a mirrored volume set is mounted, it is possible that one or more mirrored partners may be missing or not responding. The steps to be followed if a disk fails during mounting are as follows:

1. If LDEV 32's partner (LDEV 33) did not mount, LDEV 32 is placed in the PENDING state and a console message is displayed to alert you of this condition. If the partner of LDEV 32 comes online, it is recognized; otherwise, the message displays every thirty seconds.

2. Reply to the (Y/N)? question to stop it from repeating.
   This reply does nothing but stop the message from appearing on the console.

```
?09:09/12/MIRRORED PARTNER MISSING FOR LDEV# 32
?09:09/22/ACKNOWLEDGE MIRRORED PARTNER MISSING FOR LDEV#32 (Y/N)?
:REPLY 22,Y
```

3. You will **not** be able to access MEMBER2 due to its PENDING state until you do one of the following:
   a. Power on the missing partner (if the drive had previously been powered off).
   b. Issue the SUSPENDMIRRVOL command to place the PENDING disk in the SUSPEND-MIRR state and make it accessible without mirroring.

4. Use the DSTAT command to verify that LDEV 32's partner did not mount.

```
LDEV-TYPE     STATUS        VOLUME     (VOLUME SET - GEN)
30- 079370    MASTER-MD     MEMBER1    (PROD_SET-0)
31- 079370    MASTER-MD     MEMBER1    (PROD_SET-0)
32- 079370   *PENDING-MD    MEMBER2    (PROD_SET-0)
```

The previous screen shows that MEMBER2 is waiting (PENDING) for the mount of its partner (LDEV 33 which is not even listed because it failed to mount). The MEMBER2 volume remains in the PENDING state and **remains unavailable until you issue the SUSPENDMIRRVOL** command to override and tell the system to proceed without mirroring on that volume.

5. Use the VOLUTIL SUSPENDMIRRVOL command to access the MEMBER2 volume without mirroring.

NOTE The SUSPENDMIRRVOL command can only be issued on a disk in the PENDING state.

*Verify: SUSPEND THE MIRROR PENDING VOLUME ON LDEV 32 [Y/N]?**Y**

```
:VOLUTIL

  Mirvutil A.00.00, (C) Hewlett-Packard Co., 1990.
  All Rights Reserved.

  volutil:SUSPENDMIRRVOL PROD_SET:MEMBER2 32
```

CAUTION Care must be taken when using SUSPENDMIRRVOL to ensure that the PENDING disk is good. This command **forces** the system to mount and use this drive. Because of drive errors, it may not have been possible to mark the drive as bad. The drive could contain data that has not been updated. This could lead to application errors and force a reload of the volume set.

6. After you have suspended a mirrored volume, use the SHOWSET command with the MIRROR parameter to verify that the volume can be accessed and is in the SUSPEND-MIRR state.

   The previous screen shows that the MEMBER2 volume is available and does not have a mirrored partner.

```
volutil:SHOWSET PROD_SET MIRROR

Volume Name   Vol Status   Mirr Status     Ldev   Mirr ldev
MEMBER1       MASTER       NORMAL           30      31
MEMBER1       MASTER       NORMAL           31      30
MEMBER2       MEMBER       SUSPEND-MIRR     32      *
```

7. Check the disk that did not mount (ldev 33) to see if it was powered on. If it is powered off, try powering it on to see if it mounts. If it mounts, go to step 8, if not , replace the drive. Once the disk has been repaired by either physically replacing the drive or fixing the disk problem, power on the disk.

8. Use the DSTAT command to verify that the new volume can be initialized (SCRATCH or UNKNOWN status). If you need to scratch the volume and it does not contain any data that you want to save, use the SCRATCHVOL command.

NOTE : The new volume must be mounted in the SCRATCH or UNKNOWN state. It does not need to have the same LDEV or I/O path as the disk that did not mount.

```
volutil: :DSTAT

LDEV-TYPE     STATUS        VOLUME      (VOLUME SET - GEN)
30- 079370    MASTER-MD     MEMBER1     (PROD_SET-0)
31- 079370    MASTER-MD     MEMBER1     (PROD_SET-0)
32- 079370    MEMBER-MD     MEMBER2     (PROD_SET-0)
33- 079370    SCRATCH
```

NOTE: If the new volume mounts in the PENDING state, do not issue the SUSPENDMIRRVOL command on the new volume. Disk mirroring cannot work properly if both partners of a mirrored pair are placed in the SUSPEND-MIRR state.

9. Use the REPLACEMIRRVOL command to initialize LDEV 33 as the new mirrored disk partner of LDEV 32.

```
volutil: REPLACEMIRRVOL PROD_SET:MEMBER2 33
```

The system now recognizes (mounts) the replaced volume, starts the repair process, and resumes disk mirroring.

10. Use the SHOWSET command to verify that MEMBER2 is under repair.

```
volutil: SHOWSET PROD_SET MIRROR

Volume Name   Vol Status    Mirr Status     Ldev    Mirr ldev
MEMBER1       MASTER        NORMAL           30       31
MEMBER1       MASTER        NORMAL           31       30
MEMBER2       MEMBER        REPAIR-SRCE      32       33
MEMBER2       MEMBER        REPAIR-DEST      33       32
```

The above screen shows that LDEV 33 (REPAIR-DEST) is being repaired by LDEV 32 (REPAIR-SRCE). Programs and data residing on MEMBER2 are available while repairs are taking place.

☐ **User volume disk failure after system startup, power on, or mounting (other than the master disk of the user volume set)**

The system automatically recovers from a failure of a single disk that is a partner of a mirrored pair during normal mirrored operation. Normal mirrored operation means that both partners are fully mounted and no repair operation is taking place.

A drive can fail and be marked DISABLED in the following ways:

Errors being returned The drive is marked as having failed (DISABLED) immediately, and the application continues to use the remaining drive in the NON-MIRROR state.

Drive not responding There is a slight delay (40 to 60 secs) while the system waits for the drive to respond. During this waiting period, processes performing I/O will be suspended. If the drive responds before the timeout, normal mirroring resumes. If the drive does not respond, the drive is marked as having failed (DISABLED), and the application continues to use the remaining drive in the NON-MIRROR state.

To recover from a disk that failed after mounting, do the following:

1. If LDEV 32 fails, a console message alerts you of this condition.

```
?09:09/12/MIRRORED VOLUME DISABLED ON LDEV# 32
```

The system automatically continues the application without mirroring and places the good disk (ldev 33) in the NON-MIRROR state. The following message displays every thirty seconds, asking you to acknowledge this condition.

2. Your reply stops the repeating message on the console. It causes no other action to take place.

```
?09:09/22/ACKNOWLEDGE MIRRORED VOLUME DISABLED ON LDEV# 32 [Y/N]?
:REPLY 22,Y
```

NOTE: The repeating message continues until a reply is given - even if the drive is replaced.

3. Use the DSTAT command to show that LDEV 32 has failed and is no longer available.

```
:DSTAT

LDEV-TYPE     STATUS          VOLUME   (VOLUME SET - GEN)
30- 079370    MASTER-MD       MEMBER1  (PROD_SET-0)
31- 079370    MASTER-MD       MEMBER1  (PROD_SET-0)
32- 079370   *DISABLED-MD     MEMBER2  (PROD_SET-0)
33- 079370    MEMBER-MD       MEMBER2  (PROD_SET-0)
```

4. Use the SHOWSET command to confirm that LDEV 32 is disabled and that LDEV 33 is functioning in a NON-MIRROR state.

```
:VOLUTIL

Mirvutil A.00.00, (C) Hewlett-Packard Co., 1990.
All Rights Reserved.

volutil:SHOWSET PROD_SET MIRROR

Volume Name   Vol Status   Mirr Status   Ldev   Mirr ldev
MEMBER1       MASTER       NORMAL        30       31
MEMBER1       MASTER       NORMAL        31       30
MEMBER2       MEMBER       DISABLED      32       33
MEMBER2       MEMBER       NON-MIRROR    33       32
```

5. Now schedule time for the CE to replace the DISABLED disk (LDEV 32). Usually, this will have to be done during off-hours (night or weekend). Remember that you haven't lost any data but now LDEV 33 is running without protection. If it should fail before LDEV 32 is replaced, you will have to reload the volume set. Before the CE can replace the failed drive, you need to quiesce all applications that run on all volumes sets that are attached to the chain upon which the failed drive is attached. This may be as simple as getting all users to log off the application and executing a VSCLOSE volume_set_name; NOW command. Another alternative is to get all users off and shut down MPE/iX completely. You will notice that with SCSI II or SCSI Fast/Wide drives, the CE will have to power down all the drives in the SCSI chain, thus affecting any volume sets with drives on this chain. Be sure to consult your Mirrored Disk diagram before the CE begins so that you know which drives willbe affected. Also, realize that all drives that are powered down will go into repair state when powered back up. This means that if you have five drives in the SCSI chain with the one that has to be replaced, all five will go into repair state when they are powered back up (unless they were VSCLOSE'd first).

   Once the CE is done replacing and testing the new drive, it is essential that you follow this procedure before the system is restarted. This will ensure that during the mounting process, the good drive will mount first.

   a) Power the replaced disk drive down again (this may require that you power down the "hotel" in which the drive is located
   b) Restart MPE/iX with START NORECOVERY
   c) Execute "DSTAT ALL" to verify that the good drives have mounted in the PENDING State
   d) Use volutil to suspend mirroring on the replaced disks' partner.

      volutil: **SUSPEND MIRRVOL PROD_SET: MEMBER2 33**

   Use SHOWSET to verify that Ldev 33 is now in the SUSPEND_MIRR state.

```
Volume Name   Vol Status   Mirr Status   Ldev   Mirr ldev
MEMBER1       MASTER       NORMAL        30       31
MEMBER1       MASTER       NORMAL        31       30
MEMBER2       MEMBER       PENDING-MD    32       33
MEMBER2       MEMBER       PENDING-MD    33       32
```

e) Power up the replaced disk (or its hotel). The disk should mount with one of LOANER, ERROR, or DISABLED states. Use DSTAT ALL to verify this.

NOTE When replacing a volume in the DISABLED state, you must use the same LDEV number and I/O path as the failed disk. Any volume mounted on that LDEV mounts in the DISABLED state and is available as the target of the REPLACEMIRRVOL command.

6. Use the REPLACEMIRRVOL command to replace LDEV 32 and resume mirroring (after the repair).

```
volutil: REPLACEMIRRVOL PROD_SET:MEMBER2 32
```

The system now recognizes (mounts) the replaced volume starts the repair process, and resumes disk mirroring. The replacement volume has the same characteristics specified when the disabled volume was first initialized using the NEWMIRRVOL or NEWMIRRSET commands.

NOTE REPLACEMIRRVOL always initiates an immediate repair (no staging), even if the maximum number of repairs is already taking place.

If the destination drive fails or is not responding during a repair operation, it is marked DISABLED, and the source drive returns to the NON-MIRROR state.

7. Use the SHOWSET command to verify that MEMBER2 is under repair.

```
:VOLUTIL

Mirvutil A.00.00, (C) Hewlett-Packard Co., 1990.
All Rights Reserved.

volutil:SHOWSET PROD_SET MIRROR

Volume Name    Vol Status    Mirr Status    Ldev    Mirr ldev
MEMBER1        MASTER        NORMAL          30       31
MEMBER1        MASTER        NORMAL          31       30
MEMBER2        MEMBER        REPAIR-SRCE     32       33
MEMBER2        MEMBER        REPAIR-DEST     33       32
```

3.0 SUMMARY

Thorough and careful planning is critical to the successful implementation and management of Mirrored Disk/iX. If after reviewing this document along with the "Volume Management User Manual" and "Mirrored Disk/iX User's Guide", you and your staff feel unsure about the implementation of Mirrored Disk/iX, Hewlett-Packard Technical Consultants are available to assist you in these activities.

Please contact your local Hewlett-Packard Sales Office to arrange for HP Technical Consultant assistance.