

Replacing an AIX Fixed Disk

Overview

This document describes how to replace regular, striped (Raid 0) and mirrored (Raid 1) disk drives in AIX.

The procedure to replace an AIX disk depends on the situation. Considerations include disk configuration, the type of volume group and logical volumes, and the extent of the failure. Figure 1 in the appendix is a high level flowchart that helps identify the procedure to use to replace the disk. The general disk replacement steps are:

1. Identify all LV's residing on the defective disk
2. Remove all physical partitions associated with the defective disk.
3. Remove the disk from its Volume Group and from AIX
4. Replace physical disk
5. Recreate LV's
6. Restore data

Standard Disclaimer

Please use this information with care. IBM will not be responsible for damages of any kind resulting from its use. The use of this information is the sole responsibility of the customer and depends on the customer's ability to evaluate and integrate this information into the customer's operational environment.

Replacing an AIX Fixed Disk

Table of Contents

Overview	1
Standard Disclaimer	1
Table of Contents	2
Replacing an AIX Fixed Disk	3
Before a Problem Occurs	3
Special Case: Volume Group with One Disk	3
Make a Backup	3
Identify All LV's Residing on the Defective Disk	4
Remove All Physical Partitions on the Defective Disk	4
Considerations for Removing System LV's (JFSLOG, Paging, Sysdump)	5
Remove the Disk Definitions	6
Special Considerations	6
Physically Remove Disk	6
Adding a New Drive to an Existing Volume Group	6
Recreate Logical Volumes and File systems	7
Special Considerations	7
Appendix	9
Document the Disk Configuration	9
Disk Replacement Strategies	12
Figure 1: Flow Chart for Replacing Disks	13

Replacing an AIX Fixed Disk

Replacing an AIX Fixed Disk

Before a Problem Occurs

Document the disk configuration and make regular backups. Rebuilding a failed disk requires knowledge of how it was configured. It may not be possible to retrieve the configuration after a failure. For more information, see "*Documenting the LV Configuration*" and "Making a Backup" in the Appendix.

Special Case: Volume Group with One Disk

If the defective drive is the only disk in the Volume Group, then use the three step procedure:

1. **exportvg <Vgname>**
2. **rmdev -dl <hdisk#>**
3. **restore from backup**

Make a Backup

There are two general types of backup in AIX: one creates an install tape for the operating system and the other saves data.

If you are having a problem with an operating system disk (*rootvg* volume group), you need to reinstall your system. This requires a system backup (install image), which requires special backup commands such as *mksysb*, *alt_disk_install*, or equivalent. The *mksysb* command (*smitty mksysb*) is standard method for making a system backup to tape or CDROM. You can also use the *alt_disk_install* method to dynamically copy the operating system to another disk, then reboot on the new disk. These commands are limited to backing up the operating system (*rootvg*). They can not be used to backup data in *non-rootvg* volume groups. Both commands are part of the base AIX operating system

There are a number of commands to save data outside *rootvg*. They include tar, cpio, pax, backup, dd, savevg. In addition, many applications and databases have their own backup commands. Be sure to backup raw partitions as well as file systems. See the man pages for more information.

The backup frequency depends on value of the data and its volatility. The more valuable or volatile the data, the more it should be backed up. In addition, the backup quality should be verified. Tapes can become defective or the backup procedure may have a problem. Verify the backup by restoring to a different system, or by running a "table of contents." Use a group of backup tapes, and use a different tape for each backup. If a tape fails, you can use the previous backup tape.

Replacing an AIX Fixed Disk

Identify All LV's Residing on the Defective Disk

To list the LV's residing on a disk, run:

```
lspv -l <hdisk#>  
-or-  
lspv -l <PVID>
```

Sometimes the **hdisk#** becomes invalid. If so, try substituting the PVID for **hdisk#** (from the *lspv* command).

Remove All Physical Partitions on the Defective Disk

All physical partitions (PP) must be removed from the disk to be replaced. The way to remove the partitions depends on whether the disk is mirrored.

The first step is to stop applications and users from using the disk. If "normal" commands don't work, use the **fuser** command as a last resort to kill all processes using the logical volume.

```
fuser -kxuc /dev/<LVname>
```

On mirrored disks, remove partitions associated with each logical volume using the *rmlvcopy* command. The command syntax is:

```
rmlvcopy <LVname> 1 <hdisk#>
```

<hdisk#> refers to the name of the failed disk. If the *hdisk#* is unusable, use the PVID. The *<LVname>* is the logical volume name. Run this command for each LV on the failed disk. In some cases the *unmirrorvg* command may be used in place of *rmlvcopy*. See the man pages for each command to determine which command best fits your situation.

On unmirrored disks, the entire LV must be removed, even if just one of its physical partitions is on the disk to be replaced. The command syntax to remove a logical volume is

```
rmlv <LVname>
```

If the LV contains a JFS file system, unmount the file system first. Then use the *rmfs* command to remove both filesystem and logical volume at the same time.

Replacing an AIX Fixed Disk

```
umount /<FSname>  
rmfs /<fsname>
```

Removing Special LV's

Page Space: Determine if an unmirrored paging space is active.

```
lsps -a
```

If it is active, set it to be inactive on the next reboot.

```
chps -a n <LVname>  
Reboot to deactivate page space
```

Remove the inactive page space:

```
rmpps <LVname>
```

JFS Log: If the unmirrored disk contains the JFSLOG (typically named **loglvnn**), you won't be able to write to any files system in the Volume Group. To remove the JFSLOG, unmount **all** file systems in the Volume Group, then use the **rmlv** command to remove it. Record the name of the JFSLOG because you'll want to use the same name when you recreate it.

System Dump Space: If the LV is serving as a dump device, the dump pointer must first be reassigned. The same is true if the LV was mirrored and the copy is being removed. Check the dump pointers by entering:

```
sysdumpdev -l
```

Reassign the dump pointers. Enter:

```
sysdumpdev -Pp /dev/sysdumpnull (for the primary device)  
sysdumpdev -Ps /dev/sysdumpnull (for the secondary device)
```

Disk is Defective, But Not Failed: If the disk has not failed and there is another disk available, you can minimize the repair work with **migratepv** command. This command can be used to dynamically move the LVs off the failing disk to another. See man pages for more information on this command.

Replacing an AIX Fixed Disk

Remove the Disk Definitions

Removing the disk definition is a two step process. The disk must first be removed from its Volume Group, then remove the disk physical definition from AIX. Assuming we want to remove **hdisk10** from the **datavg** Volume Group, the commands would be:

```
reducevg -f datavg hdisk10  
rmdev -dl hdisk10
```

If the **hdisk** definition is no longer valid, try replacing with its PVID (found using the **lspv** command). You must also ensure that the PVID is removed from the ODM with the following command. The 32-digit value supplied consists of the PVID plus 16 zeros. For example:

```
odmdelete -q value=0073659c2c6d26f10000000000000000 -o CuAt
```

Special Considerations

SSA Disk: SSA disks also have a **pdisk** definition, which must be removed. The association between **hdisk** and **pdisk** number is random. See "smit ssa" to determine the **pdisk** number associated with the **hdisk**:

```
rmdev -dl pdisknn
```

Physically Remove Disk

To physically remove the hard disk, consult the documentation for that device, or the hardware service organization for the vendor.

In general, "hot swap" and SSA drives may be replaced while the other disks on the bus are active. Otherwise, do not attempt to remove a disk while other disks are active on the same bus. Removing a disk can cause the SCSI adapter to reset, causing operating system misdiagnose all disks on the bus as failed.

Adding a New Drive to an Existing Volume Group

If the failure was an unmirrored **rootvg** disk, reinstall the operating system on the new disk. This is normally done from a "mkysyb" image (tape or CD) or AIX install CD's. See the documentation for installing AIX for more information.

Replacing an AIX Fixed Disk

Otherwise run the **cfgmgr** command to identify the new disk to AIX. This step is not necessary if the system has been rebooted, as **cfgmgr** is run as part of the boot process.

Once the new drive has been identified, ensure that a proper PVID has been written to the drive by running:

```
chdev -l <hdisk#> -a pv=clear
chdev -l <hdisk#> -a pv=yes
```

NOTE: On SSA drives, the first chdev command may be omitted.

Add the drive to the volume group:

```
extendvg VGname hdisk#
```

If this was the only drive in the Volume Group, use the **mkvg** command to recreate the volume group on the new drive.

Recreate Logical Volumes and File systems

New logical volumes, paging spaces, file systems, or logical volume copies can be re-added with one or more of the following with the **mklv**, **mkps**, **crfs**, **mklvcopy**, **mirrorvg**, **restorevg** commands. All of these commands can be run through "smit".

If the failed disk was unmirrored, create new LVs using the **mklv** command. If you have the map file, you can use the "-m" flag to recreate the partitions in the exact order and physical location on the disk (see section on *Documenting the Disk Configuration*). If the LV contained a file system, use "smit crfs" to "Add a Journaled File Systems on a Previously Defined Logical Volume".

If the failed disk was mirrored, use the "**mklvcopy**" command. Again, you can use a MapFile to restore the copies to the exact location as before. After recreating the LV mirror copies, be sure to synchronize them using the "**syncvg**" command. For example, the "**syncvg -P 6 -v <VGname>**" synchronizes the entire volume group. See the man pages for other options..

Special Considerations

Boot Disk: If the failed disk was a mirrored *rootvg* disk, be sure it is bootable. Verify hd5 is properly mirrored on the new disk. If one does not exist, see the bosboot documentation for creating one. Define the disk as bootable to AIX. For example, if hdisk0 and hdisk1 are mirrored disks in *rootvg*, define them as bootable via the command.

Replacing an AIX Fixed Disk

```
bootlist -m normal hdisk0 hdisk1
```

JFSLOG: This is a special partition for logging JFS file system changes. There is at least one JFSLOG in every Volume Group containing a JFS file system. To recreate the JFSLOG, first unmount all JFS file systems in its volume group. Delete and recreate the JFSLOG on a disk in the same volume group using the same LV name and type JFSLOG. Format the partition. Mount the file systems.

For example, assume the log file name is **loglv01** and the volume group name is **datavg**, and the target disk is **hdisk10**. The commands to recreate the JFSLOG would be:

```
# unmount all JFS file systems datavg  
rmlv /dev/loglv01  
mklv -y loglv01 -t JFSLOG datavg hdisk10  
logform loglv01  
# remount file systems in datavg
```

Bruce Spencer

IBM

baspence@us.ibm.com

Replacing an AIX Fixed Disk

Appendix

Document the Disk Configuration

I generally use these commands to document the disk configuration:

1. `lspv`
2. `lsvg -l <VGname>`
3. `lspv -l <hdisk#>`
4. `lslv <LVname>`
5. `lslv -m <LVname>`

1 Use the `lspv` command to identify the Volume Group containing the defective disk. The output will list the hdisk, its PVID and the volume group to which it belongs:

```
# lspv
      hdisk0 00000036960cbdd1 rootvg
      hdisk1 00000036960c31de rootvg
      hdisk2 00000036960d3007 vg01
      hdisk3 000003870001328f vg01
      hdisk4 00000360ebf34660 vg01
      hdisk5 00000360d7c1f19f vg01
      hdisk6 00000036628b9724 vg02
```

2. Use the `lsvg` command to list the Logical Volumes in the Volume Group. For example, the output for `vg01` might look like:

```
# lsvg -l vg01
```

```
vg01:
LV NAME  TYPE    LPs  PPs  PVs    LV STATE    MOUNT POINT
loglv01  JFSLOG  1    1    1      open/syncd  N/A
lv04     jfs     31   31   1      open/syncd  /vol/x11r4
lv27     jfs     53   53   1      open/syncd  /vol/X11R5
lv08     jfs     13   26   1      open/syncd  /vol/tex
lv10     jfs     63   63   1      open/syncd  /vol/abc
lv21     jfs     84   84   2      open/syncd  /var/spool/news
lv12     jfs     99   99   1      open/syncd  /vol/lpp/3005
lv23     jfs     66   66   2      open/syncd  /vol/src
lv07     paging 92   92   1      open/syncd  N/A
```

This output tells us several things. First, it tells whether the logical volumes are mirrored. In this case, only `lv98` is mirrored (PP's = 2 x LP's). Second, it tells us the size of the LV's if we have to recreate them (LP's). Third, it shows that `lv21` and `lv23` reside on 2 disks (PV=2).

3. Use the `lspv -l` command to list the LV's that reside on the defective disk.

Replacing an AIX Fixed Disk

```
# lspv -l hdisk3
```

```
hdisk3:
LV NAME  LPs  PPs  DISTRIBUTION      MOUNT POINT
lv08     26   26   02..10..10..04..00 /vol/tex
lv21     16   16   02..00..12..00..02 /var/spool/news
lv04     31   31   01..08..05..15..02 /vol/x11r4
```

If the hdisk name no longer exists, and the disk is identifiable only by its 16-digit PVID, substitute the PVID for the disk name (see *lspv*). For example, hdisk3 would be

```
# lspv -l 000003870001328f
```

If you receive the following error:

```
0516-320 : Physical volume 000003870001328f is not assigned to a volume group.
```

Run the following command:

```
putlvodm -p `getlvodm -v <VGname>` <PVID>
```

For hdisk3 in vg01 type:

```
putlvodm -p `getlvodm -v vg01` 000003870001328f
```

4. The *lslv* command to lists configuration information for each LV on the defective disk. Check for non-default settings in case they need to be recreated.

```
# lslv lv04
```

```
LOGICAL VOLUME:      lv04                VOLUME GROUP:      rootvg
LV IDENTIFIER:      0000059293ce0140.9  PERMISSION:        read/write
VG STATE:           active/complete    LV STATE:          opened/syncd
TYPE:               jfs                WRITE VERIFY:      off
MAX LPs:           512                PP SIZE:           4 megabyte(s)
COPIES:            1                  SCHED POLICY:     parallel
LPs:               31                  PPs:              31
STALE PPs:         0                  BB POLICY:         relocatable
INTER-POLICY:      minimum             RELOCATABLE:      yes
INTRA-POLICY:      middle              UPPER BOUND:      32
MOUNT POINT:       /vol/x11r4          LABEL:             /vol/x11r4
MIRROR WRITE CONSISTENCY: on
EACH LP COPY ON A SEPARATE PV ?: yes
```

5. Optionally, use the "*lslv -m <LVname>*" command to map the physical location of each partition for each LV. Convert the output to a Map File format, and use it with the "*mklv -m MapFile*" or "*mklvcopy -m MapFile*" to recreate the partitions in the exact order and physical location. For example:

Replacing an AIX Fixed Disk

```
#lslv -m lv04  
LP      PP1  PV1                PP2  PV2                PP3  PV3  
0001    0016 hdisk2  
0002    0017 hdisk2  
0003    0018 hdisk2  
0004    0019 hdisk2  
0005    0020 hdisk2  
0006    0021 hdisk2  
0007    0022 hdisk2  
0008    0023 hdisk2
```

The Map File format is "PV:PP". The following "awk" command is one way of converting the "lslv -m" output to the Map File format. The example creates a Map File for the first LV copy (columns 2 and 3).

```
# lslv -m lv04 | awk ' { print $3 : $2 } '  
PV1:PP1  
hdisk2:0016  
hdisk2:0017  
hdisk2:0018  
hdisk2:0019  
hdisk2:0020  
hdisk2:0021  
hdisk2:0022  
hdisk2:0023
```

To create a map file for the second or third copy, simply change the column numbers in the "awk" print statement

Replacing an AIX Fixed Disk

Disk Replacement Strategies

Table 1: Disk Replacement Strategies

Volume Group	Unmirrored Disk	Mirrored Disk
rootvg	Backup: mksysb Procedure: <ul style="list-style-type: none"> • reinstall AIX 	Backup: mksysb Procedure: <ul style="list-style-type: none"> • Unmirror LVs residing on failed disk • Remove/replace disk • Recreate mirrors
Non-rootvg	Backup: <ul style="list-style-type: none"> • tar, cpio, backup, pax, etc Procedure: <ul style="list-style-type: none"> • Remove all PP's on failed disk • Remove/replace disk • Recreate LV's • Restore from backup 	Backup: <ul style="list-style-type: none"> • tar, cpio, backup, pax, etc Procedure: <ul style="list-style-type: none"> • Remove all LV copies on failed disk. • Remove/replace disk • Recreate LV copies

Replacing an AIX Fixed Disk

Figure 1:
Flowchart for Replacing Regular, Raid 0 and 1 Disks in AIX

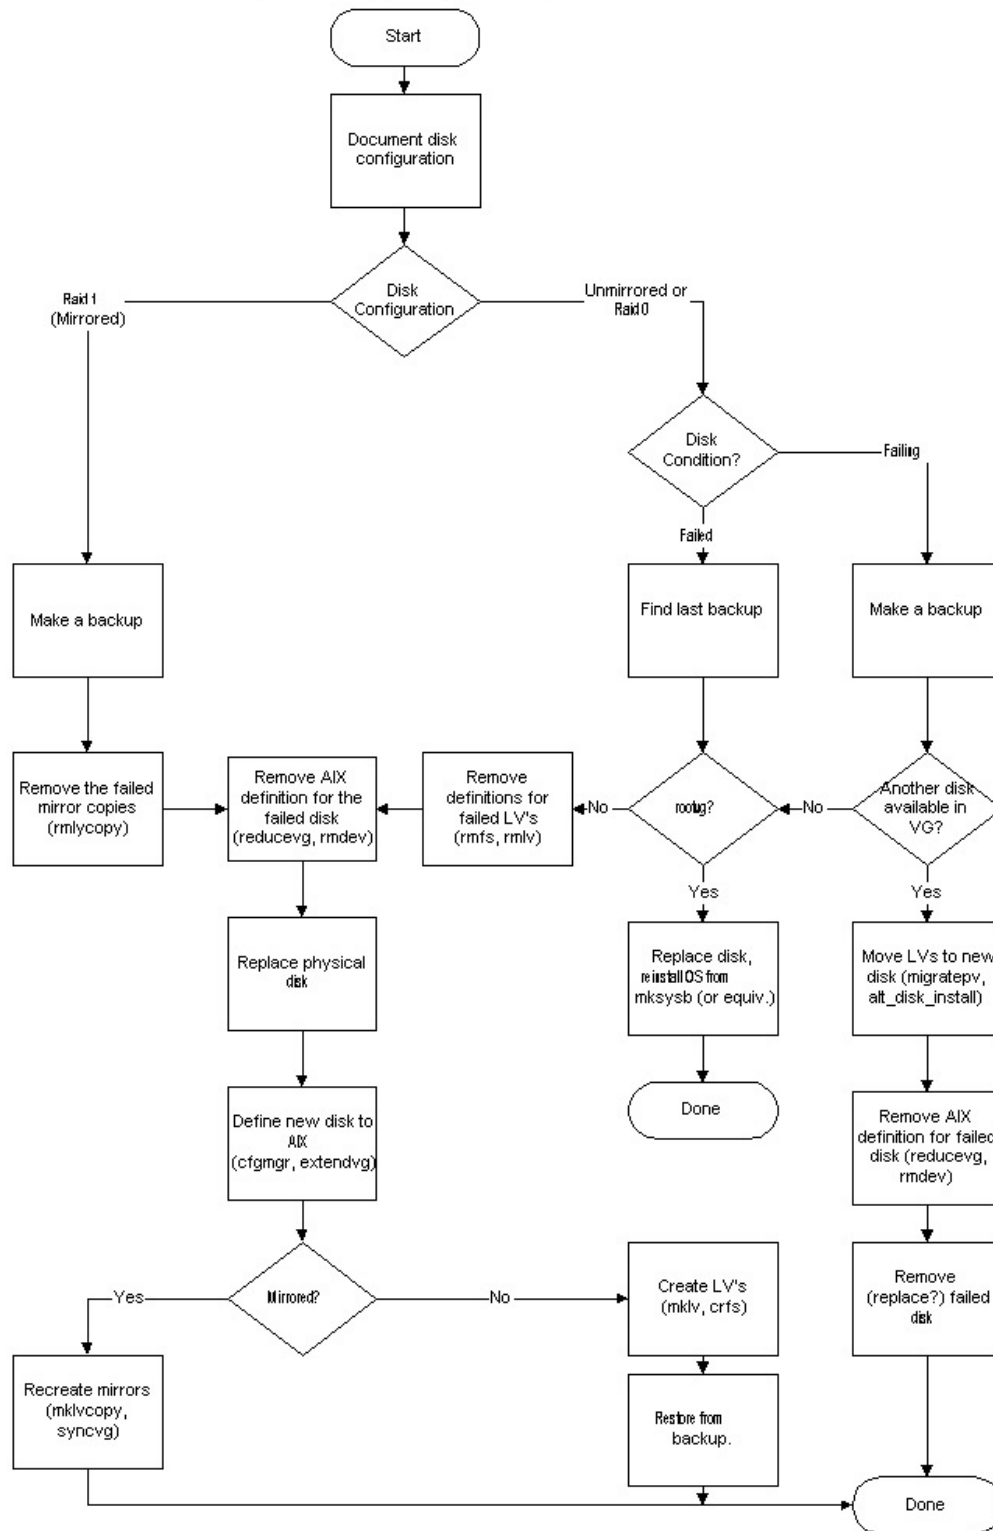


Figure 1: Disk Replacement Flowchart