

Introduction to AIX Mirror Pools

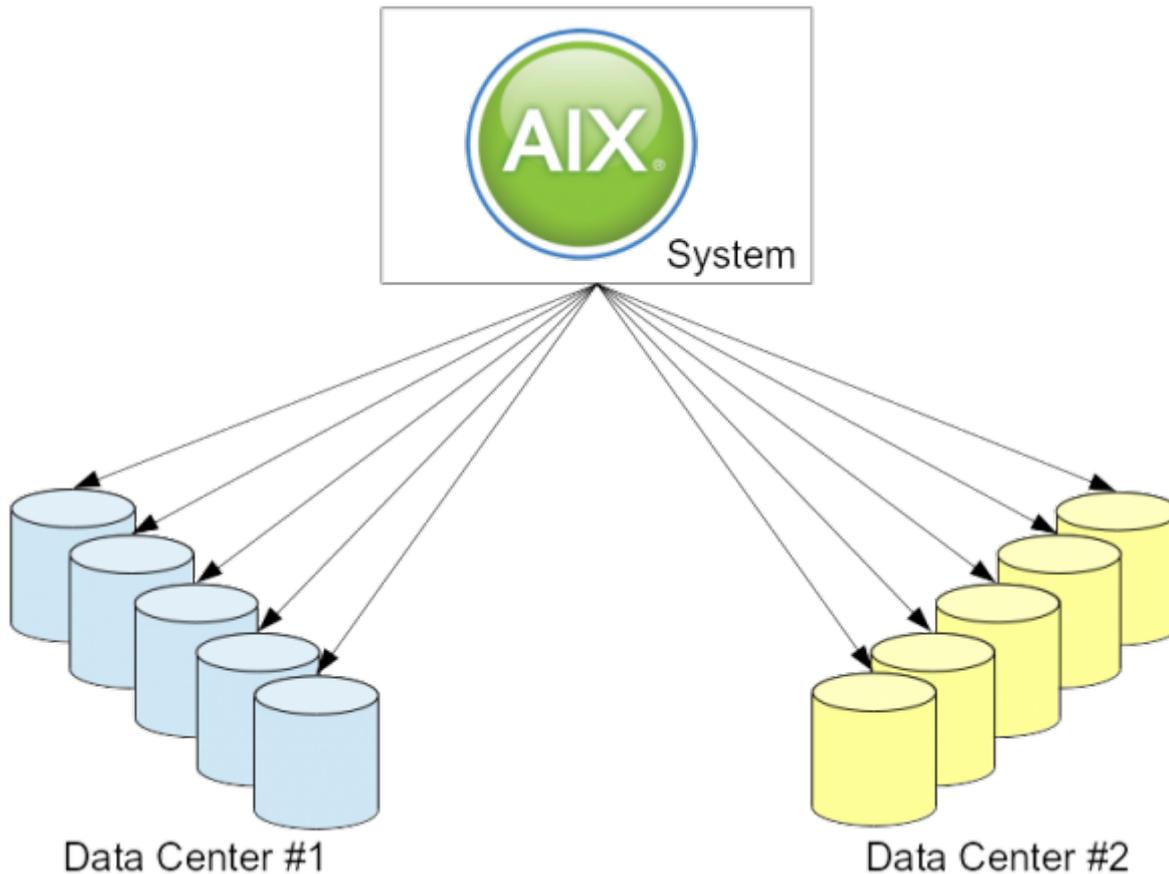
1 Introduction

Starting with AIX V6.1 Technology Level 2 so-called mirror pools were introduced that make it possible to divide the physical volumes of a scalable volume group into separate pools. The main advantage of mirror pools is that they simplify the task of mirroring data significantly compared to the steps that were required before with previous AIX versions. They allow to group physical volumes in a scalable volume group so that a mirror copy of a logical volume can be restricted to only allocate partitions from physical volumes in a specified group.

A mirror pool is made up of one or more physical volumes. Each physical volume can only belong to one mirror pool at a time. When creating a logical volume, each copy of the logical volume being created can be assigned to a mirror pool. Logical volume copies that are assigned to a mirror pool will only allocate partitions from the physical volumes in that mirror pool. This provides the ability to restrict the disks that a logical volume copy can use. Allocation of physical partitions for a logical volume can thus be restricted to a certain pool on a per copy basis.

Without mirror pools, the only way to restrict which physical volume is used for allocation when creating or extending a logical volume is to use a map file which can be very tedious. Thus, using mirror pools greatly simplifies this process.

An example of how mirror pools can be beneficial is when used with remote disks (please see drawing below). If a volume group is created with physical volumes that are located in two different locations, the disks in one location can be assigned to one mirror pool and the disks in the other location to a different mirror pool. When a logical volume is created in that volume group, each mirror copy of that logical volume can be assigned to a mirror pool. Thus, when partitions are allocated for that copy they will only come from disks that are in the assigned mirror pool.



No additional commands for mirror pools have been added to AIX but instead the existing AIX LVM commands have been extended to incorporate the mirror pool functionality. For example, mirror pools can be created with the `extendvg` command or the `chpv` command.

2 System requirements

The following system requirements must be fulfilled for mirror pools:

- Mirror pools are only available in AIX V6.1 TL 2 and higher.
- Mirror pools are only available for SVG type (scalable) volume groups.
- After assigning PVs (physical volumes) to a mirror pool, the volume group can no longer be imported to a previous version of AIX that does not support mirror pools.
- While it is possible to assign multiple logical volume copies to a mirror pool, it is recommended that only one copy of a logical volume be assigned to a mirror pool.
- Volume groups can enable strict mirror pools. If this is enabled all of the logical volumes in the volume group **must** use mirror pools.
- Any changes to mirror pool characteristics will not affect partitions allocated before the changes were made. The `reorgvg` command should be used after mirror pool changes are made to move the allocated partitions to conform to the mirror pool restrictions.

3 Technical Overview

Mirror pools is the concept of grouping disks into pools with the intent to allow logical volume mirror copies to be allocated from, or restricted to, a particular disk pool. Mirror pools allow users to restrict

the allocation of logical volume copies to a group of disks:

- Each disk in a volume group can be assigned to a mirror pool.
- Each copy of a logical volume can also be assigned to a mirror pool.
- When disk partitions are allocated to a logical volume, each copy of the logical volume will only allocate disk partitions from the disks that belong to the same mirror pool as the logical volume copy.



Mirror pools simplify the task of isolating a logical volume copy to a specific group of disks allowing for better data disaster protection.

One example of how mirror pools would be used is when an administrator wants to increase the size of a file system. Without mirror pools in order to ensure that partitions allocated to extend the file system only come from certain disks, a map file must be used. Then if he did not allocate enough new partitions for the new file system size the `extendlv` command could still be called and allocate partitions where the administrator did not want them. So the administrator needs to examine the changes made and may need to undo then redo them if they are not appropriate. Mirror pools allow the administrator to call `chfs` without having to worry about partitions being allocated from disks he does not want to be used. This kind of allocation is very useful if both local and remote disks are part of a volume group.

3.1 Creating a mirror pool

Mirror pools are created by simply assigning physical volumes to them. This is done with the `chpv` command. If a mirror pool does not exist when a physical volume is assigned to it, the mirror pool is created. Otherwise, the physical volume is simply assigned to the existing mirror pool. Physical volumes can also be assigned to mirror pools when a volume group is extended with the `extendvg` command.

Mirror pools are designated by the user via an alphanumeric name. This name can be up to 15 characters long and is unique to the volume group it belongs to. Therefore, two separate volume groups could use the same name for their mirror pools. Mirror pool names follow the same rules that apply to volume group names and logical volume names.

The `chpv` command can also be used to remove a physical volume from a mirror pool. When a mirror pool no longer has any physical volumes assigned to it, it is deleted and the mirror pool index that was being used is freed.

Mirror pools are created by assigning a physical volume to it. A mirror pool must be created before a logical volume copy can be assigned to it. Physical volumes can be assigned to a mirror pool using one of the following commands:

```
$ mkvg -S -p <mirror_pool_name> <hdisk_list>
$ extendvg -p <mirror_pool_name> <hdisk_list>
$ chpv -p <mirror_pool_name> <hdisk_list>
```

Multiple mirror pools can be created with additional `chpv -p` commands.

3.2 Using mirror pools

When creating a logical volume, mirror pools can be turned on with the `mklv` command which specifies the mirror pool to use for each copy. The mirror pools can be specified for new copies with the `mklvcopy` and `mirrorvg` commands. In addition, mirror pools can be assigned to a copy before it is created with any of the previously mentioned commands or with the `chlv` command.

When using mirror pools on logical volumes, different names will need to be passed for each copy being created. The names are a string in the form `copyN=MirrorPool` where `N` is the copy that is being specified (1-3), and `MirrorPool` is the name of the mirror pool.

For example, to create a logical volume with 3 copies, assigned to mirror pools `MPoolA`, `MPoolB`, and `MPoolC`, the following command would be used:

```
$ mklv -c 3 -p copy1=MPoolA -p copy2=MPoolB -p copy3=MPoolC <vg_name>
<lv_size>
```

Each copy of a logical volume can be assigned to a mirror pool:

```
$ mklv -p copy1=MPoolA -p copy2=MPoolB -c 2 <vg name> <lv_size>
$ mklvcopy -p copy3=MPoolC <lv_name> <number_of_copies>
$ mirrorvg -p copy2=MPoolB -p copy3=MPoolC -c 3 <vg_name>
```



Partitions allocated to a logical volume copy after the mirror pool has been assigned will only be allocated from physical volumes in the mirror pool.

3.3 Changing mirror pools

With the `chlv` command one can change to which mirror pool a copy of a logical volume is assigned to. However, any partitions allocated before the command is run will belong to the old mirror pool. The `reorgvg` command will need to be run to move the old partitions to the new mirror pool. The `reorgvg` command will also need to be run after a physical volume is moved out of a mirror pool in order to remove partitions from that mirror pool that are allocated on the physical volume.

The mirror pool a logical volume copy is assigned to can be changed as follows:

```
$ chlv -m copy1=new_mirror_pool <lv_name>
```

Physical volumes can be removed from a mirror pool:

```
$ chpv -P <hdisk_list>
```

A mirror pool can be renamed:

```
$ chpv -m <new_mirror_pool_name> <hdisk_list>
```

3.4 Viewing mirror pool information

The `lsvg`, `lspv`, and `lslv` commands can be used to view information about mirror pools. Here is a list of typical commands:

Display the PVs (physical volumes) in a volume group and their mirror pools:

```
$ lsvg -P <vg_name>
```

Display the LVs (logical volumes) in a volume group and their mirror pool:

```
$ lsvg -m <vg_name>
```

Display the mirror pool the PV (physical volume) belongs to by default:

```
$ lspv <pv_name>
```

Display the volume groups and mirror pools that all PVs (physical volumes) belong to:

```
$ lspv -P
```

Display the mirror pools that the LV (logical volume) copies belong to by default:

```
$ lslv <lv_name>
```

3.5 Strict mirror pools

Strict mirror pools is an option that can be enabled or disabled on a per volume group basis. When strict mirror pools are enabled any logical volume created in the volume group must have mirror pools enabled for each copy of the logical volume. Logical volumes created before strict mirror pools are enabled do not have to have mirror pools enabled. Strict mirror pools will be turned on/off with the `chvg -M` command or can be enabled when creating a volume group with the `mkvg -M` command. Strict mirror pool status can be viewed with the `lsvg` command.

Volume groups with strict mirror pools enabled can only create logical volume copies that have a mirror pool assigned:

```
$ mkvg -M y -S <hdisk_list>  
$ chvg -M y <vg_name>
```

A super strict allocation policy can be set so that the partitions allocated for one mirror cannot share a physical volume with the partitions from another mirror. Volume groups with super strict mirror pools enabled have three additional restrictions:

1. There can be a maximum of three mirror pools per volume group.
2. Each mirror pool must contain at least one copy of each logical volume.
3. Local physical volumes and GLVM remote physical volumes cannot belong to the same mirror pool.

Volume groups with super strict mirror pools are enabled with the same commands as strict mirror pools:

```
$ mkvg -M s -S <hdisk_list>  
$ chvg -M s <vg_name>
```

Mirror pool strictness can be disabled:

```
$ chvg -M n <vg_name>
```



- Super strict mirror pools are a requirement for using GLVM asynchronous mirroring.
- Super strict mirror pools ensure that each logical volume will have a copy on the local disks as well as on the remote disks. This ensures that if a site is lost, a full copy of the data will be available.

4 Example: How to use mirror pools

We consider the following environment where we have two separate data centers with hdisks (storage subsystems) from each center. Once mirror pools are implemented the task of extending mirrored file systems consists only of two steps:

1. Assign the new hdisks to the volume group while specifying the correct mirror pool name.
2. Extend the file system simply with chfs.

AIX then ensures the correct mirroring and strict separation of hdisks into the two separate data centers.



Especially in combination with “Poorman-Striping” (PP-Striping) AIX mirror pools drastically simplify the task of extending file systems!
PP-Striping denotes the setting of the “*LV interphysical volume allocation policy*” to “*maximum number of physical volumes*”, i.e., usage of “-e x” during `mklv`.

For the names of the mirror pools we simply use the serial number of the storage subsystem (in this example two Hitachi USP V systems) in each data center, for instance:

```
Center1-MP001  
Center2-MP002
```

The storage subsystem identifiers were obtained via the `lscfg` command (or you could ask your SAN staff):

```
hdisk6 U9117.MMA.100ED40-V3-C6-T1-W50060E3010000127-L5000000000000 Hitachi  
MPIO Disk USP V (Fibre)  
hdisk7 U9117.MMA.100ED40-V3-C6-T1-W50060E3010000237-L10000000000000 Hitachi
```

MPIO Disk USP V (Fibre)

4.1 Create a new mirror pool volume group

First, we have to find out which hdisk comes from which data center and storage subsystem:

```
$ for i in $(lsdev -Cc disk -F name); do lscfg -vL $i; done | grep hdisk
hdisk2 U9117.MMA.100ED40-V3-C6-T1-W50060E3010000127-L1000000000000 Hitachi
MPIO Disk USP V (Fibre)
hdisk3 U9117.MMA.100ED40-V3-C6-T1-W50060E3010000127-L2000000000000 Hitachi
MPIO Disk USP V (Fibre)
hdisk4 U9117.MMA.100ED40-V3-C6-T1-W50060E3010000127-L3000000000000 Hitachi
MPIO Disk USP V (Fibre)
hdisk5 U9117.MMA.100ED40-V3-C6-T1-W50060E3010000127-L4000000000000 Hitachi
MPIO Disk USP V (Fibre)
hdisk6 U9117.MMA.100ED40-V3-C6-T1-W50060E3010000127-L5000000000000 Hitachi
MPIO Disk USP V (Fibre)
hdisk7 U9117.MMA.100ED40-V3-C6-T1-W50060E3010000237-L1000000000000 Hitachi
MPIO Disk USP V (Fibre)
hdisk8 U9117.MMA.100ED40-V3-C6-T1-W50060E3010000237-L2000000000000 Hitachi
MPIO Disk USP V (Fibre)
hdisk9 U9117.MMA.100ED40-V3-C6-T1-W50060E3010000237-L3000000000000 Hitachi
MPIO Disk USP V (Fibre)
hdisk10 U9117.MMA.100ED40-V3-C6-T1-W50060E3010000237-L4000000000000 Hitachi
MPIO Disk USP V (Fibre)
hdisk11 U9117.MMA.100ED40-V3-C6-T1-W50060E3010000237-L5000000000000 Hitachi
MPIO Disk USP V (Fibre)
```

From the output above one can see that hdisk2-hdisk6 are located in data center #1 and hdisk7-hdisk11 in data center #2.

Create the volume group:

```
$ mkvg -M s -S -y testvg -s 16 -P 2048 -p Center1-MP001 hdisk2 hdisk3
$ varyonvg testvg
$ extendvg -p Center2-MP002 testvg hdisk7 hdisk8
$ lsvg -P testvg
```

Physical Volume	Mirror Pool
hdisk2	Center1-MP001
hdisk3	Center1-MP001
hdisk7	Center2-MP002
hdisk8	Center2-MP002

If you try now to create a logical volume as usual it will fail as you did not specify the mirror pools:

```
$ mklv -c 2 -t jfs2 -y lvtest1 -e x testvg 1
0516-1829 mklv: Every mirror pool must contain a copy of the logical volume.
0516-822 mklv: Unable to create logical volume.
```

The correct syntax is:

```
$ mklv -c 2 -p copy1=Center1-MP001 -p copy2=Center2-MP002 -t jfs2 -y lvtest1
-e x testvg 1
```

How do we find out where the logical volume is located?

```
$ lsvg -m testvg
Logical Volume    Copy 1          Copy 2          Copy 3
lvtest1           Center1-MP001  Center2-MP002  None
```

Finally create the file system:

```
$ crfs -v jfs2 -A yes -p rw -a logname='INLINE' -m /test1 -d lvtest1
$ mount /test1
$ df -m /test1
$ chlv -x 4096 lvtest1
$ chfs -a size=64G /test1
```

and as you can verify with

```
$ lslv -m lvtest1
```

you will see that the logical volume mirror and Poorman-Striping (PP-Striping) are in perfect shape.



Without mirror pools one would have to use map files instead which is a rather tedious and error-prone process.

4.2 Convert existing volume groups and logical volumes to mirror pools

What steps are necessary to convert an existing volume group and all of its logical volumes to mirror pools

Example: Assume the volume group and logical volumes were created without mirror pools:

```
$ mkvg -S -y testvg -s 16 -P 2048 hdisk2 hdisk3 hdisk4 hdisk5 hdisk7 hdisk6
hdisk8 hdisk9 hdisk10 hdisk11
$ varyonvg testvg
$ mklv -t jfs2 -y lvtest1 -e x testvg 1
$ crfs -v jfs2 -A yes -p rw -a logname='INLINE' -m /test1 -d lvtest1
$ mount /test1
$ chlv -x 2560 lvtest1
$ chfs -a size=40G /test1
```

Now, how can this “mess” (visible via `lslv -m lvtest1`) be corrected so we have properly mirrored logical volumes?

```
$ chvg -M s testvg
$ chpv -p Center1-MP001 hdisk2 hdisk3 hdisk4 hdisk5 hdisk6
$ chpv -p Center2-MP002 hdisk7 hdisk8 hdisk9 hdisk10 hdisk11
$ chlvs -m copy1=Center1-MP001 -m copy2=Center2-MP002 lvtest1
```

In case you receive the following message:

```
0516-1812 lchangelv: Warning, existing allocation violates mirror pools.
Consider reorganizing the logical volume to bring it into
compliance.</color>
```

then this step is still required:

```
$ reorgvg testvg lvtest1
```

Once the reorgvg process - which can take quite some time (depending on the size of the logical volume(s)) - is finished you have perfectly mirrored logical volumes with AIX mirror pools.

5 Summary

AIX mirror pools which were introduced with AIX V6.1 Technology Level 2 simplify the task of mirroring data significantly. For that purpose the existing AIX LVM commands have been extended to incorporate the mirror pool functionality. Mirror pools allow to group physical volumes in a scalable volume group so that a mirror copy of a logical volume can be restricted to only allocate partitions from physical volumes in a specified group.

Without mirror pools, the only way to restrict which physical volume is used for allocation when creating or extending a logical volume is to use a map file. This typically is a very tedious and error-prone process. Thus, the main advantage of mirror pools is that they simplify the task of mirroring data significantly compared to the steps that were required before. This is especially beneficial when used with remote disks. If a volume group is created with physical volumes that are located in two different locations, the disks in one location can be assigned to one mirror pool and the disks in the other location to a different mirror pool. When a logical volume is created in that volume group, each mirror copy of that logical volume can be assigned to a mirror pool. Thus, when partitions are allocated for that copy they will only come from disks that are in the assigned mirror pool.

6 References

- [1] [AIX 7.1 documentation mirror pools](#)
- [2] [AIX 7.1 documentation "chpv" command](#)
- [3] [AIX 7.1 documentation "mklv" command](#)
- [4] [AIX 7.1 documentation "mkvg" command](#)

From:
<https://www.perzl.org/blog/> - Michael Perzl's blog...

Permanent link:
https://www.perzl.org/blog/doku.php?id=introduction_to_aix_mirror_pools

Last update: **2018/12/19 23:15**

