

AIX rootvg Mirroring

Article Number: 18 | Rating: Unrated | Last Updated: Mon, May 21, 2018 9:30 PM

Volume group Mirroring:-

Volume group mirroring is the most essential which keeps the data into redundant way in-order to avoid possible data loss in case any hardware or disk failure. Moreover, mirroring technology has been called as RAID -1 level. We need two different disk to mirror the volume group, normally we used to have the same RAID controller disk in a partition, but considering the cost factor and other dependency we are forced to keep the same RAID controller disk for mirroring which is limitation. But, in case of that particular hardware RAID controller fail we will land in to the situation such that we will not able to access the volume groups whatever the disk come from that controller. If we keep the two different RAID controller disk which will eliminated in case of the RAID controller fault , however keeping two raid controller may not be the good option hence we will not able to accommodate may partitions in the same physical frame.

There are some principles for mirroring

- 1) Only logical volume can be mirrored direct physical disk cannot be mirrored
- 2) Each logical volume can be mirrored by two copies or maximum three copies.

In logical volume mirror mapping mechanism plays major role, each logical partition holds one physical partition and each physical partition will be resides in physical volume. To keep the track about the physical partition and logical partitions it has the unique identification number.

You can use the `lsvg -l` command to validate the mirror, please find the screen shot for the same.

The fact is that the allocation or mapping location need not be the same for mirrored disk for each logical partition to physical partition. You can validate through `lslv` command, you have to pick up the particular volume groups logical volume for the same.

Can be validated through `lslv -m` command.

The best practice is that keep the volume group as much as simple the reason behind is that , if you keep many number of disk and many number of logical volume you have the chance to lose the flexibility in terms of manageability also need to avoid to create a scattered physical partition across the physical volume can lead the performance issue. So that AIX LVM does not allow to keep two copy of logical volume in to the same physical volume.

Once all the logical volumes are mirrored the quorum will be disabled automatically which ensures 100% mirror otherwise we can say that is partial mirror.

Mirroring Practical Steps:-

STEP 1:- Check available free physical volume

```
[Test_Lpar:root:/:] lspv
hdisk0          00c45ee2c5f2712d          rootvg          active
hdisk1          none                      None
[Test_Lpar:root:/:]
```

STEP 2:- Free available disk will be shown as none; here we have to generate physical volume identifier for the free disk through chdev command

```
[Test_Lpar:root:/:] chdev -l hdisk1 -a pv=yes
hdisk1 changed
[Test_Lpar:root:/:]
```

STEP 3:- After chdev command you could see the PVID has been generated for hdisk1 through lspv command

```
[Test_Lpar:root:/:] lspv
hdisk0          00c45ee2c5f2712d          rootvg          active
hdisk1          00c45ee212c74efa          None
[Test_Lpar:root:/:]
```

STEP 4:- You could see the root volume group attributes through lsvg rootvg command whereas quorum enabled total PV's only one

```
[Test_Lpar:root:/:] lsvg rootvg
VOLUME GROUP:      rootvg                VG IDENTIFIER:    00c45ee200004c0000000014bc94668ec
VG STATE:          active                 PP SIZE:         512 megabyte(s)
VG PERMISSION:     read/write             TOTAL PPs:       558 (285696 megabytes)
MAX LVs:           256                    FREE PPs:        535 (273920 megabytes)
LVs:               14                     USED PPs:        23 (11776 megabytes)
OPEN LVs:          13                     QUORUM:          2 (Enabled)
TOTAL PVs:         1                      VG DESCRIPTORS:  2
STALE PVs:         0                      STALE PPs:       0
ACTIVE PVs:        1                      AUTO ON:         yes
MAX PPs per VG:    32512
MAX PPs per PV:    1016
LTG size (Dynamic): 256 kilobyte(s)
HOT SPARE:         no
PV RESTRICTION:    none
DISK BLOCK SIZE:   512
MAX PVs:           32
AUTO SYNC:         no
BB POLICY:         relocatable
INFINITE RETRY:    no
[Test_Lpar:root:/:]
```

STEP 5:- Now you could see hdisk0 free distribution details, which is part of rootvg.

```
[Test_Lpar:root:/:] lsvg -p rootvg
rootvg:
PV_NAME      PV STATE      TOTAL PPs   FREE PPs   FREE DISTRIBUTION
hdisk0       active        558         535       111..110..94..111..109
[Test_Lpar:root:/:]
```

STEP 6:- Check, PP's and LP's mapping, below given snap stats that LP's are equal to PP's which means this is not been mirrored

```
[Test_Lpar:root:/:] lsvg -l rootvg
rootvg:
LV NAME          TYPE      LPs      PPs      PVs      LV STATE      MOUNT POINT
hd5              boot      1        1        1        closed/syncd  N/A
hd6              paging    1        1        1        open/syncd    N/A
hd8              jfs2log   1        1        1        open/syncd    N/A
hd4              jfs2      1        1        1        open/syncd    /
hd2              jfs2      9        9        1        open/syncd    /usr
hd9var           jfs2      1        1        1        open/syncd    /var
hd3              jfs2      1        1        1        open/syncd    /tmp
hd1              jfs2      1        1        1        open/syncd    /home
hd10opt          jfs2      2        2        1        open/syncd    /opt
hd11admin        jfs2      1        1        1        open/syncd    /admin
livedump         jfs2      1        1        1        open/syncd    /var/adm/ras/livedump
lv_doonce        jfs2      1        1        1        open/syncd    /opt/DoOnceAIX
lv_pridump       sysdump   1        1        1        open/syncd    N/A
lv_auditlog      jfs2      1        1        1        open/syncd    /var/log/eprise
```

STEP 7:- Check hdisk0 logical volume distribution

```
[Test_Lpar:root:/:] lspv -l hdisk0
hdisk0:
LV NAME          LPs      PPs      DISTRIBUTION      MOUNT POINT
hd5              1        1        01..00..00..00..00  N/A
hd6              1        1        00..01..00..00..00  N/A
hd10opt          2        2        00..00..02..00..00  /opt
hd3              1        1        00..00..01..00..00  /tmp
hd1              1        1        00..00..01..00..00  /home
hd2              9        9        00..00..09..00..00  /usr
hd9var           1        1        00..00..01..00..00  /var
hd8              1        1        00..00..01..00..00  N/A
hd4              1        1        00..00..01..00..00  /
hd11admin        1        1        00..00..01..00..00  /admin
livedump         1        1        00..01..00..00..00  /var/adm/ras/livedump
lv_doonce        1        1        00..00..00..00..01  /opt/DoOnceAIX
lv_pridump       1        1        00..00..00..00..01  N/A
lv_auditlog      1        1        00..00..00..00..01  /var/log/eprise
```

STEP 8:- Check is hdisk1 having any logical volumes (sure it will not have anything)

```
[Test_Lpar:root:/:] lspv -l hdisk1
[Test_Lpar:root:/:]
```

STEP 9:- Check you disk size whether which is enough to mirror or not

```
[Test_Lpar:root:/:] bootinfo -s hdisk0
286102
[Test_Lpar:root:/:] bootinfo -s hdisk1
51200
[Test_Lpar:root:/:]
```

STEP 10:- Add the hdisk1 to rootvg by using extendvg command

```
[Test_Lpar:root:/:] extendvg rootvg hdisk1
[Test_Lpar:root:/:] lspv
hdisk0          00c45ee2c5f2712d          rootvg          active
hdisk1          00c45ee212c74efa          rootvg          active
[Test_Lpar:root:/:] █
```

STEP 11:- After addition validate the distribution by lsvg -p

```
[Test_Lpar:root:/:] lsvg -p rootvg
rootvg:
PV_NAME          PV STATE          TOTAL PPs   FREE PPs   FREE DISTRIBUTION
hdisk0           active            558         535        111..110..94..111..109
hdisk1           active            99         99        20..20..19..20..20
[Test_Lpar:root:/:] █
```

STEP 12:- You could see the difference in size, but actually we need 23 PP's only check the details through lspv For hdisk0

```
[Test_Lpar:root:/:] lspv hdisk0
PHYSICAL VOLUME:  hdisk0          VOLUME GROUP:  rootvg
PV IDENTIFIER:    00c45ee2c5f2712d VG IDENTIFIER   00c45ee200004c000000014bc94668ec
PV STATE:         active
STALE PARTITIONS: 0
PP SIZE:          512 megabyte(s)
TOTAL PPs:        558 (285696 megabytes)
FREE PPs:         535 (273920 megabytes)
USED PPs:         23 (11776 megabytes)
FREE DISTRIBUTION: 111..110..94..111..109
USED DISTRIBUTION: 01..02..17..00..03
MIRROR POOL:      None
[Test_Lpar:root:/:] lspv hdisk1
```

For hdisk1

```
[Test_Lpar:root:/:] lspv hdisk1
PHYSICAL VOLUME:  hdisk1          VOLUME GROUP:  rootvg
PV IDENTIFIER:    00c45ee212c74efa VG IDENTIFIER   00c45ee200004c000000014bc94668ec
PV STATE:         active
STALE PARTITIONS: 0
PP SIZE:          512 megabyte(s)
TOTAL PPs:        99 (50688 megabytes)
FREE PPs:         99 (50688 megabytes)
USED PPs:         0 (0 megabytes)
FREE DISTRIBUTION: 20..20..19..20..20
USED DISTRIBUTION: 00..00..00..00..00
MIRROR POOL:      None
[Test_Lpar:root:/:] █
```

STEP 13:- Mirroring the root volume group , we have to use hdisk1 for the same.

```
[Test_Lpar:root:/:] mirrorvg -S rootvg hdisk1
0516-1804 chvg: The quorum change takes effect immediately.
0516-1126 mirrorvg: rootvg successfully mirrored, user should perform
                    bosboot of system to initialize boot records. Then, user must modify
                    bootlist to include: hdisk1 hdisk0.
[Test_Lpar:root:/:] █
```

Note :- In some scenarios , if rootvg size is huge which is not been satisfied the PP distribution policy we have to use mklvcopy insider of mirrorvg command , if we use also it will fail.

STEP 14:- In mirrorvg command we used -S switch for sync volume group information with ODM in background, this we can check through lsvg command. Actually LVSA & LVCB being updated while doing this.

```
[Test_Lpar:root:/:] lsvg rootvg | grep -i stale
STALE PVs:          1
STALE PPs:          14
[Test_Lpar:root:/:] lsvg rootvg | grep -i stale
STALE PVs:          1
STALE PPs:          13
[Test_Lpar:root:/:]
```

STEP 15:- You can check in detail by using lsvg command

```
[Test_Lpar:root:/:] lsvg -l rootvg
rootvg:
LV NAME      TYPE      LPs      PPs      PVs  LV STATE  MOUNT POINT
hd5          boot      1        2        2    closed/syncd  N/A
hd6          paging    1        2        2    open/syncd    N/A
hd8          jfs2log   1        2        2    open/syncd    N/A
hd4          jfs2      1        2        2    open/syncd    /
hd2          jfs2      9        18       2    open/syncd    /usr
hd9var       jfs2      1        2        2    open/syncd    /var
hd3          jfs2      1        2        2    open/syncd    /tmp
hd1          jfs2      1        2        2    open/stale    /home
hd10opt      jfs2      2        4        2    open/stale    /opt
hd11admin    jfs2      1        2        2    open/stale    /admin
livedump     jfs2      1        2        2    open/stale    /var/adm/ras/livedump
lv_doonce    jfs2      1        2        2    open/stale    /opt/DoOnceAIX
lv_pridump   sysdump   1        1        1    open/syncd    N/A
lv_auditlog  jfs2      1        2        2    open/stale    /var/log/eprise
```

STEP 16:- Depending upon rootvg size logical volume sync will take some time. Now you could see stale PP's became zero.

```
[Test_Lpar:root:/:] lsvg rootvg | grep -i stale
STALE PVs:          1
STALE PPs:          1
[Test_Lpar:root:/:] lsvg rootvg | grep -i stale
STALE PVs:          0
STALE PPs:          0
[Test_Lpar:root:/:]
```

STEP 17:- Validate in detail,

```
[Test_Lpar:root:/:] lsvg -l rootvg
rootvg:
LV NAME          TYPE      LPs      PPs      PVs      LV STATE      MOUNT POINT
hd5              boot      1         2         2      closed/syncd  N/A
hd6              paging    1         2         2      open/syncd    N/A
hd8              jfs2log   1         2         2      open/syncd    N/A
hd4              jfs2      1         2         2      open/syncd    /
hd2              jfs2      9        18         2      open/syncd    /usr
hd9var           jfs2      1         2         2      open/syncd    /var
hd3              jfs2      1         2         2      open/syncd    /tmp
hd1              jfs2      1         2         2      open/syncd    /home
hd10opt          jfs2      2         4         2      open/syncd    /opt
hd11admin        jfs2      1         2         2      open/syncd    /admin
livedump         jfs2      1         2         2      open/syncd    /var/adm/ras/livedump
lv_doonce        jfs2      1         2         2      open/syncd    /opt/DoOnceAIX
lv_pridump       sysdump   1         1         1      open/syncd    N/A
lv_auditlog      jfs2      1         2         2      open/syncd    /var/log/eprise
```

STEP 18:- Validate the quorum status, if you would be done proper mirroring automatically quorum would be disabled otherwise quorum will not be disabled. If anyone of the physical volume failure also led insufficient VGDA which will make the partition hung.

```
[Test_Lpar:root:/:] lsvg rootvg
VOLUME GROUP:      rootvg          VG IDENTIFIER:  00c45ee200004c0000000014bc94668ec
VG STATE:          active          PP SIZE:        512 megabyte(s)
VG PERMISSION:     read/write      TOTAL PPs:      657 (336384 megabytes)
MAX LVs:           256             FREE PPs:       612 (313344 megabytes)
LVs:               14              USED PPs:       45 (23040 megabytes)
OPEN LVs:          13              QUORUM:         1 (Disabled)
TOTAL PVs:         2              VG DESCRIPTORS: 3
STALE PVs:         0              STALE PPs:      0
ACTIVE PVs:        2              AUTO ON:        yes
MAX PPs per VG:    32512
MAX PPs per PV:    1016
LTG size (Dynamic): 256 kilobyte(s)
HOT SPARE:         no
PV RESTRICTION:    none
DISK BLOCK SIZE:   512
MAX PVs:           32
AUTO SYNC:         no
BB POLICY:         relocatable
INFINITE RETRY:    no
```

STEP 19:- Check the distribution

```
[Test_Lpar:root:/home/root:] lsvg -p rootvg
rootvg:
PV_NAME          PV STATE      TOTAL PPs      FREE PPs      FREE DISTRIBUTION
hdisk0           active        558            557           112..112..111..111..111
hdisk1           active        99             77           19..18..02..20..18
```

STEP 20:- Check logical volume status on both physical volume
For hdisk0

```
[Test_Lpar:root:/:] lspv -l hdisk0
hdisk0:
LV NAME      LPs      PPs      DISTRIBUTION      MOUNT POINT
hd5           1         1        01..00..00..00..00  N/A
hd6           1         1        00..01..00..00..00  N/A
hd10opt       2         2        00..00..02..00..00  /opt
hd3           1         1        00..00..01..00..00  /tmp
hd1           1         1        00..00..01..00..00  /home
hd2           9         9        00..00..09..00..00  /usr
hd9var        1         1        00..00..01..00..00  /var
hd8           1         1        00..00..01..00..00  N/A
hd4           1         1        00..00..01..00..00  /
hd11admin     1         1        00..00..01..00..00  /admin
livedump      1         1        00..01..00..00..00  /var/adm/ras/livedump
lv_doonce     1         1        00..00..00..00..01  /opt/DoOnceAIX
lv_pridump    1         1        00..00..00..00..01  N/A
lv_auditlog   1         1        00..00..00..00..01  /var/log/eprise
```

For hdisk1

```
[Test_Lpar:root:/:] lspv -l hdisk1
hdisk1:
LV NAME      LPs      PPs      DISTRIBUTION      MOUNT POINT
hd5           1         1        01..00..00..00..00  N/A
hd6           1         1        00..01..00..00..00  N/A
hd10opt       2         2        00..00..02..00..00  /opt
hd3           1         1        00..00..01..00..00  /tmp
hd1           1         1        00..00..01..00..00  /home
hd2           9         9        00..00..09..00..00  /usr
hd9var        1         1        00..00..01..00..00  /var
hd8           1         1        00..00..01..00..00  N/A
hd4           1         1        00..00..01..00..00  /
hd11admin     1         1        00..00..01..00..00  /admin
livedump      1         1        00..01..00..00..00  /var/adm/ras/livedump
lv_doonce     1         1        00..00..00..00..01  /opt/DoOnceAIX
lv_auditlog   1         1        00..00..00..00..01  /var/log/eprise
[Test_Lpar:root:/:]
```

STEP 21:- Check the bootlist , now we have only hdisk0 , we have create & set the hdisk1 also as a boot device.

```
[Test_Lpar:root:/:] bootlist -m normal -o
hdisk0 blv=hd5 pathid=0
[Test_Lpar:root:/:]
```

STEP 22:- Create bootlist by using bosboot command

```
[Test_Lpar:root:/:] bootlist -m normal hdisk0 hdisk1
[Test_Lpar:root:/:]
```

STEP 23:- Set the bootlist with hdisk1 , incase hdisk0 failed to boot from hdisk1 partition will be booted

```
[Test Lpar:root:/:] bootlist -m normal -o  
hdisk0 blv=hd5 pathid=0  
hdisk1 blv=hd5 pathid=0  
[Test_Lpar:root:/:] █
```

mirrorvg done. both disk boot-list completed.

source: <http://www.unixdcbees.com/mirrorvg/>

Posted - Mon, May 21, 2018 7:47 PM. This article has been viewed 4912 times.

Online URL: <http://kb.ictbanking.net/article.php?id=18>