

# Using the AIX splitvg command

Article Number: 597 | Rating: Unrated | Last Updated: Mon, Jun 3, 2019 2:26 PM

## Using the AIX splitvg command.

cggibbo | Mar 31 2010 | Comments (5) | Visits (30446)

1

Just the other day, I needed to use the AIX splitvg command in order to copy some data from one system to another.

I thought I'd share the experience here.

The splitvg command can split a single mirror copy of a fully mirrored volume group into a separate "snapshot" volume group.

From the man page:

*The original volume group VGname will stop using the disks that are now part of the snapshot volume group SnapVGname. Both volume groups will keep track of the writes within the volume group so that when the snapshot volume group is rejoined with the **original volume group consistent data is maintained across the rejoined mirrors copies**. Notes:*

*To split a volume group, all logical volumes in the volume group must have the target mirror copy and the mirror must exist on a disk or set of disks. Only the target mirror copy must exist on the target disk or disks.*

*The splitvg command will fail if any of the disks to be split are not active within the original volume group.*

*In the unlikely event of a system crash or loss of quorum while running this command, the joinvg command must be run to rejoin the disks back to the original volume group.*

*There is no concurrent or enhanced concurrent mode support for creating snapshot volume groups.*

*New logical volumes and file system mount points will be created in the snapshot volume group.*

*The splitvg command is **not supported for the rootvg**.*

*The splitvg command is **not supported** for a volume group that has an **active paging space**.*

*When the splitvg command targets a concurrent-capable volume group which is varied on in non-concurrent mode, the new volume group that is created will not be varied on when the splitvg command completes. The new volume group must be varied on manually.*

So, looking at point 4, above, if you are using **enhanced concurrent volume groups** (for example with PowerHA), then you will not be able to use the splitvg command. This is disappointing as this would have been very handy in some of the large PowerHA systems I have worked with.....perhaps this will be supported in the future?

Anyway, back to my example. I had wanted to break-off one of the mirrors of a mirrored volume group and then assign the “split” volume group to another host to copy some data off.

The volume group datavg contained two disks, hdisk0 and hdisk3, as shown in the lspv output below.

```
# lspv
hdisk1      00c01c705bdc6136      old_rootvg
hdisk0      00c01c70c050810f      datavg      active
hdisk2      00c01c7018a47201      rootvg      active
hdisk3      00c01c70fed9e41a      datavg      active
```

There were only two logical volumes (loglv00, jfs2 log and fslv00, data) and a single file system (/data) in this volume group.

```
# lsvg -l datavg
datavg:
LV NAME      TYPE    LPs  PPs  PVs  LV STATE  MOUNT POINT
loglv00      jfs2log  1    2    open/syncd  N/A
fslv00      jfs2    16   32    open/syncd  /data
```

The volume group datavg, was mirrored across hdisk0 and hdisk3, as shown in the lsvg output below.

```
# lsvg -p datavg
datavg:
PV_NAME      PV STATE    TOTAL PPs  FREE PPs  FREE DISTRIBUTION
hdisk0      active      583        566      117..100..116..116..117
hdisk3      active      583        566      117..100..116..116..117
```

The logical volumes for the /data file system and the JFS2 log were both mirrored, as shown in the lslv/lspv output.

# lslv fslv00

LOGICAL VOLUME: fslv00 VOLUME GROUP: datavg  
LV IDENTIFIER: 00c01c7000004c0000000122e9038480.2 PERMISSION: read/write  
VG STATE: active/complete LV STATE: opened/syncd  
TYPE: jfs2 WRITE VERIFY: off  
MAX LPs: 512 PP SIZE: 256 megabyte(s)  
**COPIES: 2** SCHED POLICY: parallel  
LPs: 16 PPs: 32  
STALE PPs: 0 BB POLICY: relocatable  
INTER-POLICY: minimum RELOCATABLE: yes  
INTRA-POLICY: middle UPPER BOUND: 32  
MOUNT POINT: /data LABEL: /data  
MIRROR WRITE CONSISTENCY: on/ACTIVE  
EACH LP COPY ON A SEPARATE PV ?: yes  
Serialize IO ?: NO

# lslv loglv00

LOGICAL VOLUME: loglv00 VOLUME GROUP: datavg  
LV IDENTIFIER: 00c01c7000004c0000000122e9038480.1 PERMISSION: read/write  
VG STATE: active/complete LV STATE: opened/syncd  
TYPE: jfs2log WRITE VERIFY: off  
MAX LPs: 512 PP SIZE: 256 megabyte(s)  
**COPIES: 2** SCHED POLICY: parallel  
LPs: 1 PPs: 2  
STALE PPs: 0 BB POLICY: relocatable  
INTER-POLICY: minimum RELOCATABLE: yes  
INTRA-POLICY: middle UPPER BOUND: 32  
MOUNT POINT: N/A LABEL: None  
MIRROR WRITE CONSISTENCY: on/ACTIVE  
EACH LP COPY ON A SEPARATE PV ?: yes  
Serialize IO ?: NO

# lspv -l **hdisk0**

hdisk0:

LV NAME	LPs	PPs	DISTRIBUTION	MOUNT POINT
fslv00	16	16	00..16..00..00..00	/data
loglv00	1	1	00..01..00..00..00	N/A

```
# lspv -l hdisk3
```

```
hdisk3:
```

LV NAME	LPs	PPs	DISTRIBUTION	MOUNT POINT
fslv00	16	16	00..16..00..00..00	/data
loglv00	1	1	00..01..00..00..00	N/A

Using the splitvg command I was able to break-off one of the disks from the mirrored pair. This created a new volume group on hdisk3 called vg00.

```
# splitvg -c1 datavg
```

```
# lspv
```

hdisk1	00c01c705bdc6136	old_rootvg	
hdisk0	00c01c70c050810f	datavg	active
hdisk2	00c01c7018a47201	rootvg	active
<b>hdisk3</b>	<b>00c01c70fed9e41a</b>	<b>vg00</b>	<b>active</b>

The new volume group contains a new logical volume (pre-fixed with fs i.e. fsfslv00) and a file system (pre-fixed with /fs i.e. /fs/data). I can mount this file system and access the data in the file system and create and/or modify files (as shown below).

```
# mount /fs/data
```

```
Replaying log for /dev/fsfslv00.
```

```
# cd /fs/data
```

```
# ls -ltr
total 112
drwxrwxrwx  2 root  system    53248 Jan 30 2009 AIX61TL2SP2
drwxr-xr-x  2 root  system    256 Aug 05 15:24 lost+found
-rw-r--r--  1 root  system      0 Sep 15 10:29 2
```

```
# touch 3
# ls -ltr
total 112
drwxrwxrwx  2 root  system    53248 Jan 30 2009 AIX61TL2SP2
drwxr-xr-x  2 root  system    256 Aug 05 15:24 lost+found
-rw-r--r--  1 root  system      0 Sep 15 10:29 2
-rw-r--r--  1 root  system      0 Sep 15 10:46 3
```

At this point I was able to export the volume group and import it on another system. I had to re-map the Virtual SCSI disk on my VIOS first.

```
# cd
# umount /fs/data

# varyoffvg vg00
# exportvg vg00

# rmdev -dl hdisk3
```

; Re-map the disk at the VIOS layer to the other AIX LPAR.

; Configure the disk on the other LPAR (with cfgmgr) and import the volume group (with importvg).

Once I was finished with vg00 on the other LPAR, I re-mapped the disk to the original AIX LPAR, configured the disk and then re-joined the disk to the datavg volume group, with the joinvg command.

; Export the VG and remove the disk on the other AIX LPAR.

; Re-map the disk at the VIOS layer to the original AIX LPAR.

# cfmgr

# lsdev -Cc | grep hdisk3

hdisk3 Available Virtual SCSI Disk Drive

# joinvg datavg

# lspv

hdisk1	00c01c705bdc6136	old_rootvg	
<b>hdisk0</b>	<b>00c01c70c050810f</b>	<b>datavg</b>	<b>active</b>
hdisk2	00c01c7018a47201	rootvg	active
<b>hdisk3</b>	<b>00c01c70fed9e41a</b>	<b>datavg</b>	<b>active</b>

The volume group is now fully mirrored again. However the partitions are stale and need to be sync'ed. I manually re-sync the volume group with the syncvg command.

# lsvg -l datavg

datavg:

LV NAME	TYPE	LPs	PPs	PVs	LV STATE	MOUNT POINT
<b>loglv00</b>	<b>jfs2log</b>	<b>1</b>	<b>2</b>	<b>2</b>	<b>open/stale</b>	<b>N/A</b>
<b>fslv00</b>	<b>jfs2</b>	<b>16</b>	<b>32</b>	<b>2</b>	<b>open/stale</b>	<b>/data</b>

# syncvg -v datavg

# lsvg -l datavg

datavg:

LV NAME	TYPE	LPs	PPs	PVs	LV STATE	MOUNT POINT
loglv00	jfs2log	1	2	2	open/ <b>syncd</b>	N/A
fslv00	jfs2	16	32	2	open/ <b>syncd</b>	/data

More information can be found here: <http://www-01.ibm.com/support/docview.wss?uid=isg3T1010934>

Posted - Mon, Jun 3, 2019 2:26 PM. This article has been viewed 4204 times.

Online URL: <http://kb.ictbanking.net/article.php?id=597>