

# How to Remove a Virtual SCSI Disk

## Table of Contents

Overview.....	1
Removing a Physical Volume from a Volume Group.....	1
1. Deallocating the physical partitions.....	1
2. Remove the physical volume from the volume group.....	3
3. How to map the virtual scsi disk (on the client partiton) to the physical disk (on the server partition)....	3
4. Remove the virtual scsi disk definition from the device configuration database on the VIO client.....	5
5. Remove the backing device on the VIO server.....	5
Related Documentation.....	6

## Overview

This document describes the procedure to remove a virtual disk in a volume group on a Virtual I/O Client, to map the virtual scsi disk to its corresponding backing device, and to remove the backing device from the Virtual I/O Server. Please, read the entire document before proceeding.

This document applies to AIX version 5.3 and above.

In a Virtual I/O environment, the physical devices are allocated to the VIO server. When there is a hardware failure (disk or adapter may go bad) on the VIO server, unless the VIO server has some type of redundancy, that will have an impact on the VIO client whose virtual disks are being served by the failing device. The impact may be loss of connectivity to the virtual scsi disks, unless there is some type of redundancy (MPIO or LVM mirroring) on the client partition.

This document does NOT apply to any of the following environments:

- If the virtual disk is in a shared volume group (i.e HACMP, etc)
- If the virtual disk is part of rootvg volume group.

## Removing a Physical Volume from a Volume Group

The following steps are needed to remove a virtual disk from the VIO client, and they are later discussed in more detail:

1. Deallocate all the physical partitions associated with the physical volume in the volume group.
2. Remove the physical volume from the volume group
3. Map the virtual scsi disk on the VIO client partiton to the backing device on the VIO server.
4. Remove the virtual scsi disk definition from the device configuration database.
5. Remove the backing device.

At this point, a new virtual scsi can be added to the VIO client in place of the virtual disk that was removed in the case where this procedure was done as a result of a hardware failure on the VIO server partition.

### 1. Deallocating the physical partitions

In the following procedure, we will be using `hdisk4` in the example, as the virtual scsi disk wanting to be removed from the VIO client.

First, we need to determine the logical volumes defined on the physical volume we want to remove. This can be done by running:

```
# lspv -l hdisk#
```

where `hdisk#` is the virtual scsi disk to be removed.

#### Example

## # **lspv -l hdisk4**

```
hdisk4:
LV NAME LPS  PPs  DISTRIBUTION      MOUNT POINT
fslv00      2    2    00..02..00..00..00  /test
loglv00 1    1    00..01..00..00..00  N/A
rawlv      30   30   00..30..00..00..00  N/A
```

If the hdisk name no longer exists, and the disk is identifiable only by its 16-digit PVID (you might see this from the output of **lsvg -p <VGname>**), substitute the PVID for the disk name. For example:

## # **lspv -l 00c2b06ef8a9f98a**

You may receive the following error:

```
0516-320 : Physical volume 00c2b06ef8a9f98a is not assigned to a volume group.
```

If so, run the following command:

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

*VGname* refers to your volume group, *PVID* refers to the 16-digit physical volume identifier, and the characters around the **getlvodm** command are grave marks, the backward single quote mark. The **lspv -l <PVID>** command should now run successfully. To determine the VGname associated with that physical volume use **lspv hdisk#**.

If another disk in the volume group has space to contain the partitions on this disk, and the virtual scsi disk to be replaced has not completely failed, the **migratepv** command may be used to move the used PPs on this disk. See the man page for the **migratepv** command on the steps to do this.

If the partitions cannot be migrated, they must be removed. The output of the **lspv -l <hdisk#>**, or **lspv -l <PVID>**, command indicates what logical volumes will be affected. Run the following command on each LV:

## # **lslv <LVname>**

The COPIES field shows if the LV is mirrored. If so, remove the failed copy with:

## # **rmlvcopy <LVname> 1 <hdisk#>**

*hdisk#* refers to all the disks in the copy that contain the failed disk. A list of drives can be specified with a space between each. Use the **lslv -m <LVname>** command to see what other disks may need to be listed in the **rmlvcopy** command. If the disk PVID was previously used with the **lspv** command, specify that PVID in the list of disks given to the **rmlvcopy** command. The **unmirrorvg** command may be used in lieu of the **rmlvcopy** command. See the man pages for **rmlvcopy** and **unmirrorvg**, for additional information.

If the logical volume is not mirrored, the entire logical volume must be removed, even if just one physical partition resides on the drive to be replaced and cannot be migrated to another disk. If the unmirrored logical volume is a JFS or JFS2 file system, unmount the file system and remove it. Enter:

## # **umount /<FSname>**

## # **rmfs /<FSname>**

If the unmirrored logical volume is a paging space, see if it is active. Enter:

## # **lspas -a**

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

## # **chps -a n <LVname>**

Then deactivate it and remove it by entering:

## # **swapoff /dev/<LVname>**

## # **rmpps <LVname>**

Remove any other unmirrored logical volume with the following command:

## # **rmlv <LVname>**

## 2. Remove the physical volume from the volume group.

In the case where the virtual scsi disk to be replaced is the only physical volume in the volume group, then remove the volume group, via:

```
# exportvg <VGname>
```

This will deallocate the physical partitions and will free up the virtual disk. Then, remove the disk definition, as noted on step 3.

In the case where there are more than one physical volumes. Using either the PVID or the hdisk name, depending on which was used when running **lspv -l** in the preceding discussion, run one of the following:

```
# reducevg <VGname> <hdisk#>
# reducevg <VGname> <PVID>
```

If you used the PVID value and if the **reducevg** command complains that the PVID is not in the device configuration database, run the following command to see if the disk was indeed successfully removed:

```
# lsvg -p <VGname>
```

If the PVID or disk is not listed at this point, then ignore the errors from the **reducevg** command.

## 3. How to map the virtual scsi disk (on the client partition) to the physical disk (on the server partition)

In the following example, we are going to determine the mapping of virtual scsi disk, hdisk4

### On the VIO client

The following command shows the location of hdisk4:

```
# lscfg -vl hdisk4
```

```
hdisk4      U9117.570.102B06E-V1-C7-T1-L810000000000    Virtual SCSI Disk Drive
```

where V1 is the LPAR ID (in this case 1), C7 is the slot# (in this case 7), and L81 is the LUN ID.

Take note of these values.

Next, determine the client SCSI adapter name, by **'grep'**ing for the location of hdisk4's parent adapter, in this case, V1-C7-T1:

```
# lscfg -v | grep V1-C7-T1
```

```
  vscsi4      U9117.570.102B06E-V1-C7-T1      Virtual SCSI Client Adapter
              Device Specific.(YL).....U9117.570.102B06E-V1-C7-T1
hdisk4      U9117.570.102B06E-V1-C7-T1-L810000000000    Virtual SCSI Disk Drive
```

where vscsi4 is the client SCSI adapter.

### On the HMC

Run the following command to obtain the LPAR name associated with the LPAR ID

```
# lshwres -r virtualio --subtype scsi -m <Managed System Name> --level lpar
```

To get the managed system name, run

```
# lssyscfg -r sys -F name
```

Then, look for the "lpar\_id" and "slot\_num" noted earlier. In our case, the VIO client lpar id is 1 and the slot # is 7.

In the following example, the managed system name is Ops-Kern-570. The VIO client partition name is kern1.

The VIO Server partition name is reg33\_test\_vios.

```
# lshwres -r virtualio --subtype scsi -m Ops-Kern-570 --level lpar
```

```
...
lpar_name=kern1,lpar_id=1,slot_num=7,state=1,is_required=0,adapter_type=client,
remote_lpar_id=11,remote_lpar_name=reg33_test_vios,remote_slot_num=23,backing_devices=
none
...
```

Take note of the remote\_lpar\_id (11) and the remote\_slot\_num (23). Then, in the same output, look for a line that corresponds to "lpar\_id 11, slot # 23"

```
...
lpar_name=reg33_test_vios,lpar_id=11,slot_num=23,state=1,is_required=0,adapter
_type=server,
remote_lpar_id=any,remote_lpar_name=,remote_slot_num=any,backing_devices=none
...
```

So in this case, VIO server reg33\_test\_vios is serving virtual scsi disk, hdisk4, on the VIO client, kern1.

### On the VIO Server

Go to the VIO Server associated with the LPAR ID obtained in the previous step, in our case reg33\_test\_vios.

As padmin, run the following command to display the mapping, which should match the mapping obtained from the HMC obtained above.

```
$ lsmmap -all|grep <VIO server lpar ID>-<VIOS slot#>
```

### Example

```
$ lsmmap -all|grep V11-C23
```

where V11 is the VIO server lpar\_id and C23 is the slot #

The cmd will return something similar to

```
vhost21      U9117.570.102B06E-V11-C23      0x00000001
```

In this case, vhost21 is the server SCSI adapter mapped to our VIO client lpar id 1 (0x00000001).

Next, list the mapping for the vhost# obtained previously.

```
$ lsmmap -vadapter vhost21
```

SVSA	Physloc	Client Partition ID
-----	-----	-----
vhost21	U9117.570.102B06E-V11-C23	0x00000001

  

```
VTD   virdisk01
LUN   0x8100000000000000
Backing device   clientlv01
Physloc
```

Take note of the VTD and Backing device name. In this case, the backing device mapped to virtual scsi disk, hdisk4, is logical volume, clientlv01, and it is associated with Virtual Target Device, virdisk01.

## **4. Remove the virtual scsi disk definition from the device configuration database on the VIO client**

To remove the vscsi definition, run

```
# rmdev -dl hdisk#
```

Ensure you know the backing device associated with the virtual scsi disk being removed prior to issuing the `rmdev` command. That information will be needed in order to do clean up on the server partition. Refer to the section "How to map the virtual scsi disk (on the client partition) to the physical disk (on the server partitions)".

## 5. Remove the backing device on the VIO server

The peripheral device types or backing devices currently supported are

- logical volume
- physical volume
- optical device starting at v1.2.0.0-FP7 (but not currently supported on System i)

Prior to removing the backing device, the virtual target device must be removed first. To do so, run the following as `padmin`:

```
$ rmdev -dev <VTD name>  
$ rmlv <LVname>
```

or you can remove both the VTD and logical volume in one command by running:

```
$ rmvdev -vtd <VTD name> -rmlv
```

In the case where the backing device is a physical volume, then, removing the virtual target device completes this document.

If you need to determine the physical device and volume group that the logical volume belongs to, you can issue the following commands prior to running `rmlv` or `rmvdev`.

```
$ lsiv -pv <LVname>
```

Lists the physical volume that the logical volume specified resides on.

```
$ lsiv <Lvname>
```

Shows the characteristics of the logical volume, including the volume group name, # of mirrored copies, and so on.

In our example, the backing device is a logical volume, `clientlv01`, and it resides on the physical device, `hdisk3`:

```
$ lsiv -pv clientlv01
```

```
clientlv01:N/A  
PV          COPIES      IN BAND      DISTRIBUTION  
hdisk3      080:000:000  100%        000:080:000:000:000
```

```
$ rmdev -dev virdisk01
```

```
virdisk01 deleted
```

```
$ rmlv clientlv01
```

```
Warning, all data contained on logical volume clientlv01 will be destroyed.
```

```
rmlv: Do you wish to continue? y(es) n(o)? y
```

```
rmlv: Logical volume clientlv01 is removed.
```

## Related Documentation

- [IBM System p Advanced POWER Virtualization Best Practices Redbook](#)
- [IBM System Hardware Information Center](#)
- [VIOS Commands Reference](#)