

HOWTO: Setup Dual VIOS VSCSI and MPIO

Contributed by Michael Felt

This article is based on a cheat sheet I use when setting up virtual scsi client partitions. Just enough detail is provided to get the job done. The testing I leave up to you!

The basis of setting up dual VIO servers will always be:

- * Setup First Virtual I/O Server
- * Test Virtual I/O setup
- * Setup Second Virtual I/O Server
- * Test Virtual I/O setup

Once the Virtual I/O Servers are prepared focus moves to the client partitions

- * Setup client partition(s).
- * Test client stability

In the following - no hints are provided on testing. Thus, not 4 parts, but two parts: setup VIOS; setup client. The setup is based on a DS4300 SAN Storage server. Your attribute and interface names may be different.

Setting up Multiple Virtual I/O Servers(VIOS)

The key issue is the LUNs exported from the SAN environment must always be available to any VIOS that is going to provide the LUN via the Virtual SCSI interface to Virtual SCSI client partitions (or client LPAR). The hdisk attribute that must be set on each VIOS for each LUN that will be used for a MPIO configuration is `reserve_setting`. When there are problems the first thing I check is whether the disk attribute `reserve_setting` has been set to none.

But it is not just this one setting. For failover we also need setup the adapters. Note: to change any of these attributes the adapters and disks need to be in a defined state - OR - you make the changes to the ODM and reboot. This second, i.e. reboot, option is required when you cannot get the disks and/or adapters offlined (varyoffvg followed by several `rmdev -l` commands).

To make an ODM change just add the parameter `-P` to the `chdev -l` commands below. And be sure to not use `rmdev -d -l <device>` as this just removes the device definition from the ODM and you have to restart.

Setup adapters

This example is using an IBM DS4000 storage system. Your adapter names may be different.

First, make sure the disks are offline. If no clients are using the disks this is easy, otherwise you must get the client partition to varyoffvg the disks, and use `rmdev -l` to each disk/LUN. Once the disks are not being used the following commands bring the disks and adapters offline (or into a defined state). The command `cfgmgr` will bring them back online with the new settings active. Do not run `cfgmgr` until AFTER the disks attributes have been set.

Note: I have used the command `oem_setup_env` to switch from user `padmin` to `root`. Below I will switch back to `padmin` (note `$` or `#` prompt for commands for `padmin` and `root`, respectively).

```
# rmdev -l dar0 -R
```

```
# rmdev -l fscsi0 -R
```

Update the adapter settings:

```
# chdev -l fscsi0 -a fc_err_recov=fast_fail
```

```
# chdev -l fscsi0 -a dyntrk=yes
```

Setup Disks

On my system the disks the first four disks are local and I have 8 LUNs that I want to make MPIO ready for the clients. So, with both VIO servers having the disks offline, or only one VIO server active, it is simple to get the disks setup for `no_reserve`. I also add that I want the VIO server to be aware of the `pvid` the clients put, or modify on these disks. Finally, I make sure the client sees the disk as "clean".

```
# for i in 4 5 6 7 8 9 10 11
```

```
> do
```

```
> chdev -l hdisk$i -a pv=yes
```

```
> chdev -l hdisk$i -a reserve_policy=no_reserve
```

```
> chpv -C hdisk$i ## this "erases the disk so be careful!
```

> done

Now I can activate the disks again using:

```
# cfgmgr
```

Setup Clients

I assume you already know how to configure VSCSI clients (you have already configured at least two VIO Servers!!) so I'll skip over the commands on the HMC for configuring the client partition.

Once the client LPAR has been configured - but before it is activated you need to map the LUNs available on the VIOS to the client. The command to do this is `mkvdev`.

Assuming that the VIOS are setup to have identical disk numbering any command you use on the first VIOS can be repeated on the second VIOS.

```
$ mkvdev -vdev hdisk4 -vadapter vhost0
```

```
vtscsi0 Available
```

You can verify the virtual target (vt) mapping using the command:

```
$ lsmap -vadapter vhost 0
```

Repeat these commands on the other VIOS(s) and make sure they are all connected to the same VSCSI client (VX-TY). The VX part is the most important - that signifies that it is the same client partition (X == client partiion number). The TY is the Virtual I/O bus slot number and this will be different for each mapping. The TY part signify the multiple paths.

The easiest way to setup the client drivers is to just install it. The AIX installation sees the MPIO capability and install the drivers and creates the vpaths. After AIX reboots there are few changes we want to make sure AIX recivers automatically incase one of the VIO servers goes offline, or loses connection with the SAN.

On the client

For each LUN coming from the SAN do:

```
# chdev -l hdiskX -a hcheck_mode=nonactive -a hcheck_interval=20 \  
-a algorithm=failover -P
```

As this is probably including your rootvg disks a reboot is required.