

OSU Tier 3 Raid Array Management

S. Durkin
Jan. 7, 2011

Overview: NFS mounted disk volumes are physically located on the following machines:

disk-0-1	/data
disk-0-0 (a.k.a. cms-0, grid-0-0)	/data1, /data2

The disks are controlled by a single Areca ARC-1280ml SATA II Raid controller on each machine. The users manual is available on <http://www.areca.us/products/pcie341.htm>. The physical map from the raid disk number to position within the Supermicro 4U case is shown below:

```
                Front View
                Top
Left Side      18 19 20 21 22 23
                12 13 14 15 16 17
                6  7  8  9 10 11
                0  1  2  3  4  5
```

For disk-0-1 disk0 is not connected since the linux boot and operating system are on that disk. For disk-0-0 disk0 and disk1 are not connected since the linux boot and operating system are on those disks.

Software: The areca disks can be managed by three methods: 1) on boot type F6 on the Areca bios prompt, 2) once booted use the Areca command line interface software, and 3) one booted use the web based Areca software.

Command Line Software: Logon to the controller computer (either disk-0-1 or grid-0-0). As super-user at the command prompt type: /usr/sbin/areca_cli. You can navigate through the commands using *command -h*. The Areca users manual is convenient to understand the commands available. If you are doing anything serious the web software is preferable.

Web Software: Logon to the controller computer (either disk-0-1 or grid-0-0). As super-user start the proxy server at the command prompt by typing: /usr/sbin/areca_httpd. On the screen you should see:

```
Copyright (c) 2004 Areca, Inc. All Rights Reserved.
Areca HTTP proxy server, Version: V1.82, Arclib: 260, Date: Sep 11 2007

Starting HTTP Proxy Server...Please wait(MAX = 5 minutes)
Controller(s) list
-----
Controller[1](PCI) : Listen to port[81].
Cfg Assistant : Listen to port[82].
Binding IP: [0.0.0.0]
Note: IP[0.0.0.0] stands for any ip bound to this host.
-----

#####
Press CTRL-C to exit program!!
#####
```

You want to talk to the port for Controller[1], port 81 chosen by the proxy server above. You then start your web server on cms-in0 entering http:10.1.255.254:81 (grid-0-0) or http: 10.1.255.251:81 (disk-0-1) (assuming the port number is 81 in the proxy server output). You will be prompted for login. The user name is admin, and the password is 0000. Follow the Areca users manual for help in navigating the menu.

On Disk Failure: An alarm will sound. All disks in the array will be red except the failed drive where no led will be illuminated. Automatic failure emailing has been set up on disk-0-1 to inform durkin and gwilliams of failure. It has not been activated yet on grid-0-0. To see what is going on with the event, logon to the controller computer (either disk-0-1 or grid-0-0). Then start the command line server by typing /usr/sbin/areca_cli. At the command prompt type:

- 1) sys beeper p=0 (turns off the beeper)
- 2) disk info (show you the disk that is in error)
- 3) event info (gives you the failure mode of the disk)

Check that the disk in the software matches the disk in the array that is in error. If not contact Stan. If so we hotswap the bad disk:

With the computer on remove the disk in error.

Get a good disk from PRB3122 in a box at the bottom of the shelves.

Replace the bad disk with the good disk. Put the bad disk in a box and leave it on Stan's desk.

With the computer on insert the good disk.

In about 10 seconds the array will start rebuilding. It can take up to 24 hours to rebuild. The disk can be accessed normally during the rebuild process.

Appendix: Building, Formatting, and Mounting a Disk Array: Insert the disks into the Supermicro slots and use one of the three utilities (preferably the web based utility so the machine remains usable) to define a Raid 6array. Follow the default values in the Areca raid array book. At this point the raid array will begin to build which may take up to 20 hours.

Once the array has been build one has to create the volume. As root:

1. fdisk /dev/sdb (make sure you have the right disk)
2. pvcreate -Z y /dev/sdb
3. vgcreate data /dev/sdb
4. lvcreate -L xxxxxx (where xxxx is in MB)
5. vgdisplay

Now mount the disk and format it (again as root):

1. `mkfs -t xfs /dev/mapper/⟨⟨volume-name⟩⟩`
(do an `ls -ltr` to find the newest volume)
2. `mkdir /mountpoint`
(where `/mountpoint` is where you want the disk mounted)
3. `mount -t xfs /dev/mapper/⟨⟨volume-name⟩⟩ /mountpoint`
4. Add the following lines to `/etc/fstab`:

```
/dev/mapper/⟨⟨volume-name⟩⟩          xfs      defaults      1 2
```