

GtL *modpod* System Reference

Thomas Naughton
Oak Ridge National Laboratory
naughtont@ornl.gov

March 17, 2004

Draft \$Id: modpod.tex,v 1.7 2004/03/17 16:11:34 tjn Exp \$

1 Introduction

This document provides details about the Genomes to Life (GtL) system called *modpod* ¹.

1.1 System Administration

The general system administration is currently managed by Thomas Naughton (naughtont@ornl.gov). George Hill (hillgw@ornl.gov) manages the Oracle database as well as the system (including database) backups. (see also Section 3, page 2)

1.2 Backup Policy

There is an overall system backup run daily. Currently this is using the Legato NetWorker client for Linux. This sends data to a central tape backup server maintained by George Hill under the condition that backup tapes will be purchased as needed ²

The system is configured with hardware RAID 5 on the three large (approx 1TB) disk arrays. This additional fault-tolerance plus the backup mechanism form a sound base for development codes as well as the scientific data stored on the system.

1.3 Recovery Policy

The Legato NetWorker recovery method is to be used. This entails the base operating system installation, network configuration and NetWorker client install/setup.

2 System Configuration

The *Relion 430* storage server is under a 3-year warranty ³ from Penguin Computing. There were some initial issues with the 3ware device driver support. Also, the default Red Hat AS 2.1 kernel was incapable of supporting > 1TB filesystems so the array configurations had to be adjusted so all three had 5 disks, yielding approximately 1TB per array unit.

The system details are:

- Serial Num: P0336084
- Receipt date: 9/30/2003

¹The actual DNS name is *modpod.csm.ornl.gov*.

²Tapes are purchased in boxes of 10 and as of the last purchase, Dec. 2003, cost approximately \$700.

³Initial purchase and receipt of system in Oct. 2003

- 3 x 3ware 7506-8 RAID cards
- 16 x Maxtor 250GB 7200 RPM IDE Hard Drives (w/ 8MB)
- 6 x Viking 1GB DDR PC2100 Memory
- 2 x Intel Xeon 3.06Ghz CPU 533Mhz FSB (w/ 512)
- Tyán S2720GN Motherboard, 533Mhz FSB
- Mitsumi 24x CD-rom
- Mitsumi Floppy
- Warranty, Penguin 3 year card
- 4U rack mountable chassis
- Intel Gigabit Ethernet 8492MT NIC (actually has 5 NICs)
 - NICs eth0-3 are Gigabit Ethernet and eth4 is FastEthernet
 - eth0 MAC 00:07:E9:1A:39:0D
 - eth1 MAC 00:07:E9:1A:39:0C
 - eth2 MAC 00:E0:81:25:7A:1F
 - eth3 MAC 00:E0:81:25:7A:1E
 - eth4 MAC 00:E0:81:25:7A:07
- Red Hat Enterprise Linux AS 2.1 Premium support edition (1st year of `up2date` service is free)

2.1 Disk Array Configuration

The three hardware RAID controllers. These have five drives attached to each controller and appear to the system as three large scsi drives. There is a 3ware tool for viewing the underlying array in `/root/bin/tw_cli`. This can be used to view the status of the array and their attached drives. For the curious, the 16th hard drive is a simple IDE drive (`/dev/hda`) and is used to house the base operating system.

When working with the 3ware CLI you are able to see the underlying details. The devices break down into three array units 0 – 2. Attached to each array unit are drives, labeled as *ports*. The current configuration has 5 drives attached to each array unit, labeled as ports 0 – 4. The diagram in Figure 1 details where these logical drives map to in the physical chassis – helpful when you have to replace a faulty drive ⁴ The information in Table 1 shows the physical to logical mapping of the array, SCSI device and filesystem.

Disk Mapping: device / filesystem

Array Unit	SCSI Dev	Disk Dev	Part/FS	Notes
u0	scsi0	sda	/dev/sda1: /u1	Verified 3/15/04
u2	scsi1	sdb	/dev/sdb1: /u2	Verified 3/17/04
u3	scsi2	sdc	/dev/sdc1: /u3	Verified 3/17/04

Table 1: The disk mappings for the RAID arrays.

See the 3ware documentation regarding the array controllers and management of *DEGRADED* arrays or for array *REBUILD/CREATION*. Note, I have found that when trying to do a rebuild from the `tw_cli` things throw an error and you may have to reboot and perform the action from the BIOS interface.

3 Contact Information

The following section provides some useful contact information. See Table 2.

⁴TJN: The scsi1 and scsi2 mappings to physical positions on the array/controller diagram (Figure 1) need to be verified.

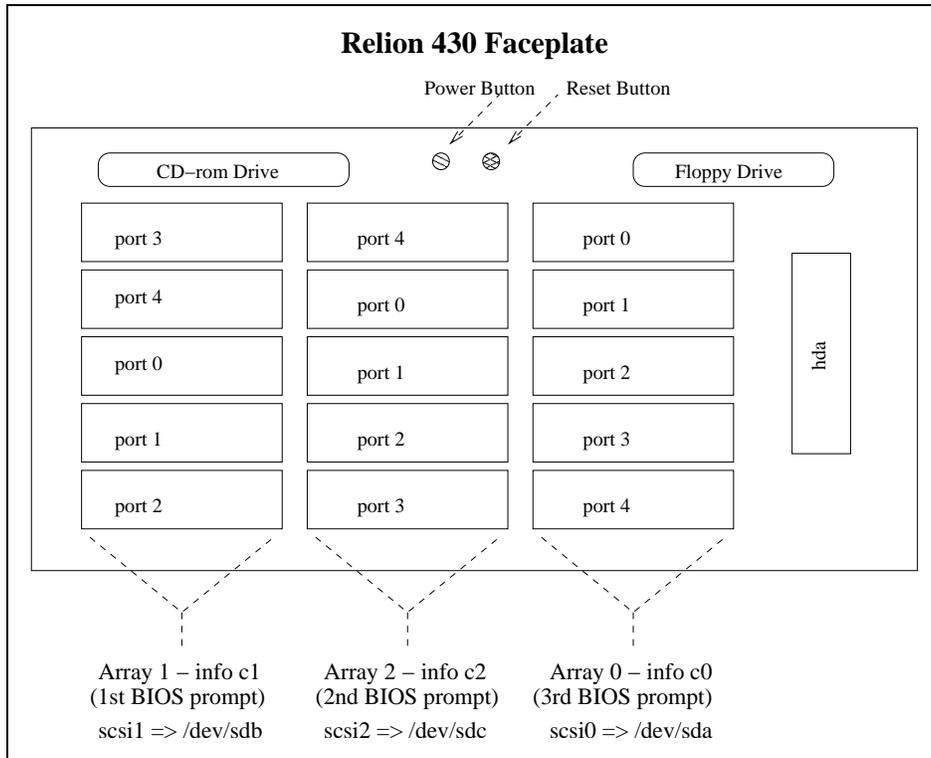


Figure 1: Diagram of the Relion 430 disk array.

Contact Information

Name	Phone Number	Email	Fax	Role
Thomas Naughton	865-576-4184	naughtont@ornl.gov	865-576-5491	General SysAdmin
George Hill	865-576-8521	hillgw@ornl.gov	-	Oracle Admin and backups
Al Geist	-	geistgaii@ornl.gov	865-576-5491	System Owner
John Mugler	865-574-9539	muglerj@ornl.gov	-	Alt. Contact
Stephen Scott	865-574-3144	scottsl@ornl.gov	-	Alt. Contact

Table 2: Contact information for the *modpod* system.

A Disk Partitioning

The filesystems on this machine are created in the normal fashion. Since the arrays appear as large SCSI devices you simply have to partition the drive and then create a filesystem. The following output shows this for the first device, scsi0 (/dev/sda).

```
[root@modpod tmp]# fdisk /dev/sda
...
# create one big primary (1) partition of type ext3 over entire device
...
[root@modpod tmp]# mke2fs -j /dev/sda1
...
[root@modpod tmp]# mount -t ext3 /dev/sda1 /u1
[root@modpod tmp]# ls -l /u1
total 16
drwx----- 2 root root 16384 Mar 15 18:34 lost+found
```