

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
HAYSTACK OBSERVATORY
WESTFORD, MASSACHUSETTS 01886

December 9, 2014

Telephone: 781-981-5951
Fax: 781-981-0590

To: Distribution
From: Chester Ruzczyk
Subject: Mark6 Maintenance

Maintenance for the Mark6 systems can be separated into two areas: the Mark6 Host / Server and the disk modules. In this memo, a description of how to maintain this equipment is presented.

1. Mark6

In the Mark6 system the major replaceable components are the mother board (and associated hardware), the OS system drive, the disk module backplane, the eSATA cables and the power supplies. The components that have the highest failure rate, based upon the manufactures estimated mean time between failures are the host chassis power supply, at 95,150 hrs, the expansion chassis power supply, 135,268 hrs, the motherboard, 500,000 hrs, and the eSATA cables, 250 insertions.

The power supplies can be monitored by observing the Mark6 general operation. If the power supplies are starting to fail, the power and current draw required to support the disk modules will usually result in a system crash.

The eSATA cables are rated for 250 insertions. It is recommended for maintenance that after the approximately 250 insertions, the cable is reversed. Reversal of the cable is to unplug the front and back of the cables from the module and SAS controller card of the Mark6 and reverse them, i.e. the cables connected to the back are now to be inserted into the front of the disk module, and the connectors of the front, are statically inserted in the SAS controller card at the back of the Mark6. Please maintain the cable with the white tape to be inserted into the top SAS slot.

Information on software maintenance and upgrade will be posted on the forum, see the contact section for details. Software utilities and test software, as developed, will be posted on the VLBI forum along with descriptions in the Mark6 User's manual.

2. Disk Modules

Disk module maintenance falls under three categories: verification or conditioning of new disk modules, operations with a failed disk within the module and, replacement of a failed disk in a module. Each of these topics is addressed below.

2.1 Conditioning of New Disk Modules

The conditioning of new disk modules can be accomplished with the software utility `hammer.sh` and the results of the conditioning process viewed with `hammerplot.sh`. The instructions for using the software utilities can be found by executing the “--help” after the command or in the Mark6 users manual.

2.2 Operations with a Failed Disk

If during operations a disk within the open group of modules fails the Mark6 software will detect the failure and remove it from the existing pool of disks that are being recorded to. There may be a slight performance degradation (in the sustainable write data rates resulting in a small percent of lost data) during recording but the system will continue normal operations.

Future versions of software will report the disk module where the faulty disk(s) exist.

2.3 Replacement of a Failed Disk

If a HDD fails in a module, the disk must first be identified. There are two methods to identify the failed HDD: 1) look at the lights on the front of the module when writing, the LED will not be illuminated if the disk is not being written, or 2) from a command line on the Mark6 issue the Linux “df” command if the disk is mounted, or the “lsscsi -t” command if it is not.

The relationship between the LEDs on the front of the module and individual HDD can only be guaranteed if the eSATA cables have been properly connected. If the cables are randomly inserted into the disk module it can be one of two disks that have failed.

If the “df” command is utilized a one to one correspondence can be made to the mount point and the device allocation. If the “lsscsi -t” command is used you will have to look at the output and determine the device that is not detected.

Once the LED or the lsscsi position is identified, the relationship between the slots and the LED's are as follows:

LED	lsscsi identifier	Printed Circuit Board Label
0	<code>sas:0x4433221100000000 / sas:0x4433221108000000</code>	B3
1	<code>sas:0x4433221101000000 / sas:0x4433221109000000</code>	B2
2	<code>sas:0x4433221102000000 / sas:0x443322110A000000</code>	B1
3	<code>sas:0x4433221103000000 / sas:0x443322110B000000</code>	B0
4	<code>sas:0x4433221104000000 / sas:0x443322110C000000</code>	A3
5	<code>sas:0x4433221105000000 / sas:0x443322110D000000</code>	A2
6	<code>sas:0x4433221106000000 / sas:0x443322110E000000</code>	A1
7	<code>sas:0x4433221107000000 / sas:0x443322110F000000</code>	A0

To remove the disk, simply loosen (do not remove) the four screws securing it to the disk module and slide it out of the module (the screws aid in guiding the HDD in and out of the module and hence should not be completely removed yet). Once the HDD is removed from the module, remove the screws and insert the screws into the new HDD. Slide the new HDD into the disk module and once the HDD connectors are mated with the female connector of the module tighten down the screws.

To format the disk for use within the module a few steps are required and must be done manually. The steps are to partition the HDD, format the HDD and copy over the appropriate meta data on to the HDD.

The steps which should be performed as root, are:

- 1) Create the disk's data partition
 - a) `parted /dev/sdX mktable gpt`
 - b) `parted /dev/sdX mkpart primary xfs 1 -100M`
 - c) `parted /dev/sdX MSN#+_disk number replacing`
 - i) e.g. `parted /dev/sdX EHT%0001_4`
- 2) Create the disk's meta-data partition
 - a) `parted /dev/sdX mkpart primary xfs -100M 100%`
 - b) `parted /dev/sdX MSN#+_disk number replacing "m"`
 - i) e.g. `parted /dev/sdX EHT%0001_4m`
- 3) Format the data / meta-data partitions with an xfs file system
 - a) `mkfs.xfs -f /dev/SDX1`
 - b) `mkfs.xfs -f /dev/SDX2`
- 4) Manually mount the disk
 - a) `mount -t xfs -o "mark6 gid" /dev/SDX2 "/mnt/disk/.meta/"slot module is in"/"disk #"`
 - i) e.g. `mount -t xfs -o 1002 /dev/sdX2 /mnt/disk/.meta/1/4`
- 5) Manually mount another disk of the module
 - a) `mount -t xfs -o "mark6 gid" /dev/SDY2 /mnt/disk/.meta/"slot module is in"/"different disk #"`
 - i) e.g. `mount -t xfs -o 1002 "/dev/SDY2 /mnt/disk/.meta/1/1`
- 6) cp over the meta data group, subgroup (if used), disk_sn and eMSN files.
 - a) `cd /mnt/disk/.meta/1/4`
 - i) `cp /mnt/disk/.meta/1/1/group .`
 - ii) `cp /mnt/disk/.meta/1/1/eMSN .`

- iii) `cp /mnt/disk/.meta/1/1/disk_sn .`
- 7) Modify the `disk_sn` to now remove the HDD replaced with the serial number of the new disk. This will have to be done for all HDD in the module.
- 8) Unmount the disks within the disk module
 - i) `cd`
 - ii) `mount -u /mnt/disk/.meta/1/1`
 - iii) `mount -u /mnt/disk/.meta/1/4`
- 9) All disks should now be available for recording.

3. Contact Information

For support of the Mark6 system it is recommended that you search the Mark6 forum <http://vlbi.org/vlbi-forums/>, and if a posting does not match your problem please make a post. The forum will contain links to existing and new Mark6 memos to help resolve maintenance issues.

If there is a hardware failure, Conduant can also be contacted via email support@conduant.com concerning the appropriate failed item. Note, that Conduant will not help you debug the potential problem, but will help find a replacement if under warranty.