

GLAST Balloon Flight Software

Subject: Disk Rescue

Version: 1.0

Created: 20 December 2001

Revised: 20 December 2001

0 General Disk Layout

The balloon flight flew six SCSI disks (numbered 0-5 by SCSI ID) with each disk segmented into nine partitions (numbered 0-8). The first eight partitions on each disk were used for event data, and the ninth partition (which was much smaller) was used for a combination of configuration data, housekeeping data and message logging.

Each partition was a mount point and the mount points followed a strict naming convention. To find a file on, for instance disk 3 partition 5, the fully qualified path name would begin with `/sd3p5`. For event data partitions, the partition name was always followed by a directory named `/evt`. The last partition held three types of data and the data types were segregated into three directories, `/cfg` for configuration data, `/hsk` for housekeeping data and `/dtt` for message logs.

File names also followed a strict naming convention. A typical (unqualified) file name would look like:

```
aaaaaaaaaa_bbbb.ccc
```

```
aaaaaaaaaa  Digit string. Timestamp (UTC) when file was opened
bbbb        Digit string. Disambiguate two timestamps less than a second apart
ccc         One of evt, cfg, hsk or dtt (you figure it out!)
```

Thus fully qualified file names on the balloon disks looked like:

```
/sd4p1/evt/0996946952_0000.evt
/sd4p8/cfg/0996848441_0000.cfg
```

1 The Disk Round Robin

To avoid exposure to a single disk failure, flight software divided the data-taking into segments and made sure that successive segments went to different disks.

Dealing with the simpler case first, the `hsk` and `dtt` files were closed and reopened every ten minutes exactly. The file changes for these two files were done in the same piece of code, and the filenames for these two files should always be “paired” (i.e. have the same value for the timestamp piece of the file name). Writing started in `/sd0p8`, move on to `/sd1p8` and so on until it reached `/sd5p8`, at which point it would roll back to `/sd0p8`.

The `evt` and `cfg` files were also “paired”, but the rollover was less deterministic. A file change could be triggered by either a time limit (ten minutes during the balloon flight) or by exceeding a file length (1 GByte for the balloon flight). The `cfg` round robin follows the `hsk/dtt` scheme. The `evt` round robin was a little more complicated because it also looped over a list of partitions. Event writing

started on /sd0p0, went through /sd1p0 and on to /sd5p0 at which point it switched to /sd0p1 and on through /sd5p1. Once it reached /sd5p7, it went back to /sd0p0.

One final subtlety. We were concerned that should writing always start on the same disk after a reboot, and should that disk fail, we'd end up in a permanent boot loop. To avoid this, the “current” writing location for both the `hsk/dtt` pair and the `evt/cfg` pair were stored in non-volatile memory and a reboot always incremented from the current write locations before opening any files. Thus when I said that writing started on /sd0p0 (or /sd0p8), I was lying. Writing actually started one disk further on from wherever the previous writing left off.

2 So How Many Disks Did We Lose?

OK. Having laid the groundwork on naming conventions and the philosophy of how we selected the disk/partition we'd write to, how did we do on writing balloon flight data, and how much of it is accessible? The following table summarizes my recovery efforts.

		Disk (SCSI ID #)					
		0	1	2	3	4	5
Partition	0	Mounts, then I/O errors	OK (recovered)	OK (recovered)	Mounts, then I/O errors	OK (recovered)	OK (recovered)
	1	Mounts, then I/O errors	OK (recovered)	OK (see notes)	Mounts, then I/O errors	OK (recovered)	OK (see notes)
	2	Mounts, then I/O errors	OK (empty)	OK (empty)	Mounts, then I/O errors	OK (empty)	OK (empty)
	3	Mounts, then I/O errors	OK (empty)	OK (empty)	Mounts, then I/O errors	OK (empty)	OK (empty)
	4	Mounts, then I/O errors	OK (empty)	OK (empty)	Mounts, then I/O errors	OK (empty)	OK (empty)
	5	Mounts, then I/O errors	OK (empty)	OK (empty)	Mounts, then I/O errors	OK (empty)	OK (empty)
	6	Mounts, then I/O errors	OK (empty)	OK (empty)	Mounts, then I/O errors	OK (empty)	OK (empty)
	7	Mounts, then I/O errors	OK (empty)	OK (empty)	Mounts, then I/O errors	OK (empty)	OK (empty)
	8	Won't mount	OK (recovered)	OK (recovered)	Won't mount	OK (recovered)	OK (recovered)

It looks like disk 0 and disk 3 are pretty much shot. Disks 1 and 4 on the other hand are operating normally. Disks 2 and 5 seem to be somewhere in between. Partitions 0 and 8 on each of disks 2 and 5 are fine, but partition 1 has some problems. In each case, an attempt to copy a file from this partition resulted in an early abort. In the case of disk 2, 80% of the file was recoverable: in the case of

disk 5, only 0.8% (there was only one file in each of these partitions). Murphy's law of course dictates that these files were real flight data!

3 Files Recovered

The following tables list all the files I recovered, segregated by type. Note that I have put the files in chronological order (so they dance around the disks). I did it this way because I find it easier to understand what files are missing this way.

Also note that these tables list *all* the files I recovered. I made no distinction on the basis that some were taken before the balloon liftoff and some after. The ones from before liftoff might be of questionable value, but they were there and available (and I have no great desire to go back to the flight disks again!)

Filename (configuration files)	Size (bytes)	File Opened (Texas time)
/sd1p8/cfg/0996839917_0000.cfg	808	2001-08-03::06:58:37
/sd2p8/cfg/0996844183_0000.cfg	808	2001-08-03::08:09:43
/sd4p8/cfg/0996845383_0000.cfg	808	2001-08-03::08:29:43
/sd5p8/cfg/0996848441_0000.cfg	808	2001-08-03::09:20:41
/sd1p8/cfg/0996852302_0000.cfg	808	2001-08-03::10:25:02
/sd2p8/cfg/0996854356_0000.cfg	808	2001-08-03::10:59:16
/sd4p8/cfg/0996857080_0000.cfg	808	2001-08-03::11:44:40
/sd5p8/cfg/0996910198_0000.cfg	808	2001-08-04::02:29:58
/sd1p8/cfg/0996911765_0000.cfg	808	2001-08-04::02:56:05
/sd2p8/cfg/0996912365_0000.cfg	808	2001-08-04::03:06:05
/sd4p8/cfg/0996913415_0000.cfg	808	2001-08-04::03:23:35
/sd5p8/cfg/0996917458_0000.cfg	808	2001-08-04::04:30:58
/sd1p8/cfg/0996920031_0000.cfg	808	2001-08-04::05:13:51
/sd2p8/cfg/0996925754_0000.cfg	808	2001-08-04::06:49:14
/sd4p8/cfg/0996926432_0000.cfg	808	2001-08-04::07:00:32
Liftoff!		
/sd5p8/cfg/0996944019_0000.cfg	808	2001-08-04::11:53:39
/sd1p8/cfg/0996945219_0000.cfg	808	2001-08-04::12:13:39
/sd2p8/cfg/0996945819_0000.cfg	808	2001-08-04::12:23:39
/sd4p8/cfg/0996946952_0000.cfg	808	2001-08-04::12:42:32
/sd5p8/cfg/0996947503_0000.cfg	808	2001-08-04::12:51:43

Filename (message logs)	Size (bytes)	File Opened (Texas time)
/sd5p8/dtt/0996798248_0000.dtt	91	2001-08-02::19:24:08
/sd1p8/dtt/0996844183_0000.dtt	3396	2001-08-03::08:09:43
/sd2p8/dtt/0996844783_0000.dtt	3170	2001-08-03::08:19:43
/sd4p8/dtt/0996848441_0000.dtt	3346	2001-08-03::09:20:41
/sd5p8/dtt/0996849041_0000.dtt	1029	2001-08-03::09:30:41
/sd1p8/dtt/0996854356_0000.dtt	337	2001-08-03::10:59:16
/sd2p8/dtt/0996856480_0000.dtt	3396	2001-08-03::11:34:40
/sd4p8/dtt/0996910198_0000.dtt	3346	2001-08-04::02:29:58
/sd5p8/dtt/0996910798_0000.dtt	111	2001-08-04::02:39:58

Filename (message logs)	Size (bytes)	File Opened (Texas time)
/sd1p8/dtt/0996912365 0000.dtt	3170	2001-08-04::03:06:05
/sd4p8/dtt/0996917458 0000.dtt	665	2001-08-04::04:30:58
/sd5p8/dtt/0996919430 0000.dtt	3346	2001-08-04::05:03:50
/sd1p8/dtt/0996925754 0000.dtt	3396	2001-08-04::06:49:14
/sd2p8/dtt/0996926354 0000.dtt	219	2001-08-04::06:59:14
Liftoff!		
/sd4p8/dtt/0996944019 0000.dtt	3450	2001-08-04::11:53:39
/sd5p8/dtt/0996944619 0000.dtt	3278	2001-08-04::12:03:39
/sd1p8/dtt/0996945819 0000.dtt	3316	2001-08-04::12:23:39
/sd2p8/dtt/0996946419 0000.dtt	3278	2001-08-04::12:33:39

Filename (event data)	Size (bytes)	File Opened (Texas time)
/sd1p0/evt/0996839917 0000.evt	9158404	2001-08-03::06:58:37
/sd2p0/evt/0996844183 0000.evt	20680544	2001-08-03::08:09:43
/sd5p0/evt/0996848441 0000.evt	21070188	2001-08-03::09:20:41
/sd1p0/evt/0996852302 0000.evt	4500612	2001-08-03::10:25:02
/sd2p0/evt/0996854356 0000.evt	397540	2001-08-03::10:59:16
/sd4p0/evt/0996857080 0000.evt	15912052	2001-08-03::11:44:40
/sd5p0/evt/0996910198 0000.evt	19859180	2001-08-04::02:29:58
/sd1p0/evt/0996911765 0000.evt	19678128	2001-08-04::02:56:05
/sd2p0/evt/0996912365 0000.evt	19387436	2001-08-04::03:06:05
/sd4p0/evt/0996913415 0000.evt	3183728	2001-08-04::03:23:35
/sd5p0/evt/0996917458 0000.evt	3290896	2001-08-04::04:30:58
/sd1p1/evt/0996920031 0000.evt	3372880	2001-08-04::05:13:51
/sd2p0/evt/0996925754 0000.evt	21183764	2001-08-04::06:49:14
/sd4p0/evt/0996926432 0000.evt	2585772	2001-08-04::07:00:32
Liftoff!		
/sd5p0/evt/0996944019 0000.evt	80184028	2001-08-04::11:53:39
/sd1p1/evt/0996945219 0000.evt	650728440	2001-08-04::12:13:39
/sd2p1/evt/0996945819 0000.evt	829708800	2001-08-04::12:23:39
/sd4p1/evt/0996946952 0000.evt	1073763116	2001-08-04::12:42:32
/sd5p1/evt/0996947503 0000.evt	7703040	2001-08-04::12:51:43

Filename (housekeeping data)	Size (bytes)	File Opened (Texas time)
/sd5p8/hsk/0996798248 0000.hsk	171152	2001-08-02::19:24:08
/sd1p8/hsk/0996844183 0000.hsk	658272	2001-08-03::08:09:43
/sd2p8/hsk/0996844783 0000.hsk	658892	2001-08-03::08:19:43
/sd4p8/hsk/0996848441 0000.hsk	658692	2001-08-03::09:20:41
/sd5p8/hsk/0996849041 0000.hsk	206968	2001-08-03::09:30:41
/sd1p8/hsk/0996854356 0000.hsk	18248	2001-08-03::10:59:16
/sd2p8/hsk/0996856480 0000.hsk	658692	2001-08-03::11:34:40
/sd4p8/hsk/0996910198 0000.hsk	658892	2001-08-04::02:29:58
/sd5p8/hsk/0996910798 0000.hsk	43352	2001-08-04::02:39:58
/sd1p8/hsk/0996912365 0000.hsk	658892	2001-08-04::03:06:05
/sd4p8/hsk/0996917458 0000.hsk	106468	2001-08-04::04:30:58
/sd5p8/hsk/0996919430 0000.hsk	658692	2001-08-04::05:03:50

Filename (housekeeping data)	Size (bytes)	File Opened (Texas time)
/sd1p8/hsk/0996925754_0000.hsk	658892	2001-08-04::06:49:14
/sd2p8/hsk/0996926354_0000.hsk	37064	2001-08-04::06:59:14
Liftoff!		
/sd4p8/hsk/0996944019_0000.hsk	658692	2001-08-04::11:53:39
/sd5p8/hsk/0996944619_0000.hsk	658472	2001-08-04::12:03:39
/sd1p8/hsk/0996945819_0000.hsk	658292	2001-08-04::12:23:39
/sd2p8/hsk/0996946419_0000.hsk	658472	2001-08-04::12:33:39

4 Where Are They Now?

I have copied all these files to the SLAC FTP space. Look for the directory...

```
/nfs/farm/g/glast/u02/bfem/data/rescue
```

... and you'll find the files I've listed above with exactly the same filenames.

Oh, and a special note for Heather. I've repaired the two files that aborted during the copy from the balloon disks. They should both pass "XIC" without complaint!

5 Where To From Here?

This is about as far as I can go. Any further recovery should be put in the hands of a professional disk recovery service (I may have already made their lives impossible by simply spinning up the disks). The decision to do that would obviously need a cost/benefit analysis. Looking at the pattern of disk usage, I could see recovering:

Filename (event data)	Size (bytes)	Comment
/sd0p1/evt/09969?????_0000.evt	~500000000	WAG
/sd2p1/evt/0996945819_0000.evt	~200000000	Remaining 20%
/sd3p1/evt/09969?????_0000.evt	~1000000000	WAG
/sd5p1/evt/0996947503_0000.evt	~1000000000	Remaining 99.3%
/sd0p2/evt/09969?????_0000.evt	~1000000000	May or may not exist

So there's probably between 2.7 and 3.7 GByte of event data still unrecovered. The "easiest" recovery would probably be /sd5p1 because the disk is basically functional (and represents good bang for the buck ... 1 GByte of data). /sd2p1 would be about the same degree of difficulty, but a lower return (only 200 MByte). Disks 0 and 3 I would probably put in the trash can!