DRAGNET Cluster Support Info

- The DRAGNET cluster was delivered on Thu+Fri 9+10 July 2015.
- We have 4 years support on the complete system.

⇒ DRAGNET cluster support ends in July 2019.

Failing Components

After initial post-delivery replacements, hard drives fail the most often (as expected). As of writing (Aug 2017), we have been through 3 replacement calls (only a few drives each time) + 1 upcoming. Other defective components can be reported via the support site or by e-mail. Please also provide some "proof" that it is defective, or discuss with them how to do that (possibly after reading their service support Tips/Troubleshooting/... pages).

For up to 20 drives, our vendor wants defective *hard drives* to be reported via their service site (instead of by e-mail). For the service we did so far, we don't need service credits (as long as our support contract lasts).
Log in, then:

- Click on one of the support call links at the bottom of the page. E.g. for hard drives:
- Verify shipping address for return parts, then fill in:
- ClusterVision machine tag: 150037 (I'm entirely unsure about this, but it doesn't matter.)
- Your machine reference: drg02 drg12 drg16
- Hard Disk Brand: TOSHIBA
- Hard Disk Model: MC04ACA400E
- Hard Disk Serialnumbers: 55I7K0XUFLSA 55IAK0N0FLSA 55IAK0NEFLSA (example)
- Number of Hard Disks with Failures: 3
- Linux Kernel Version: `uname -a` command output, e.g.: Linux drg02 3.10.0-229.el7.x86_64 #1 SMP Fri Mar 6 11:36:42 UTC 2015 x86_64 x86_64 x86_64 GNU/Linux

Hardware Settings:

1. Has the BIOS configuration been set to AHCI ?: yes
2. Is the disk part of an array on a dedicated hardware RAID controller?: no (drgXX nodes), or yes (dragnet or dragproc node)
3. Have you tried to fully re-provision the software onto the disk?: yes

Fault Description - Please use this part for all disks.:

```
[root@drg02 ~]$ smartctl --all /dev/sdb
smartctl 6.2 2013-07-26 r3841 [x86_64-linux-3.10.0-229.el7.x86_64] (local
build)
Copyright (C) 2002-13, Bruce Allen, Christian Franke, www.smartmontools.org

=== START OF INFORMATION SECTION ===
Device Model:     TOSHIBA MC04ACA400E
[...]
Error 647 occurred at disk power-on lifetime: 7064 hours (294 days + 8
hours)
```

```
    When the command that caused the error occurred, the device was active or
idle.

    After command completion occurred, registers were:
    ER ST SC SN CL CH DH
    -- -- -- -- -- -- --
    40 41 98 70 ea 28 40  Error: UNC at LBA = 0x0028ea70 = 2681456

    Commands leading to the command that caused the error were:
    CR FR SC SN CL CH DH DC   Powered_Up_Time   Command/Feature_Name
    -- -- -- -- -- -- -- --   ----------------  --------------------
    60 08 98 70 ea 28 40 00  20d+14:42:31.049  READ FPDMA QUEUED
    ef 90 03 00 00 00 a0 00  20d+14:42:31.047  SET FEATURES [Disable SATA
feature]
    ef 10 02 00 00 00 a0 00  20d+14:42:31.047  SET FEATURES [Enable SATA
feature]
    27 00 00 00 00 00 e0 00  20d+14:42:31.047  READ NATIVE MAX ADDRESS EXT
[OBS-ACS-3]
    ec 00 00 00 00 00 a0 00  20d+14:42:31.046  IDENTIFY DEVICE

Error 646 occurred at disk power-on lifetime: 7064 hours (294 days + 8
hours)
    When the command that caused the error occurred, the device was active or
idle.
[...]


[SMART info for each drive]
```

Hard Drives

The `smartd` service has been set up on all DRAGNET nodes to auto-report S.M.A.R.T. (self-monitoring for hard drives) failures (actually counters exceeding certain thresholds) by e-mail to dragnet[AT]astron[DOT]nl.
Note 1: By definition you cannot fully (or at all) count on such predictive failure analysis.
Note 2: To avoid endless spam, only 1 e-mail is sent. Another one for every new extended test failure or threshold exceeded.

To dump all S.M.A.R.T. information for a drive to a .log file, run (as root or via `sudo`):

```
$ smartctl --all /dev/XXX >> smartctl-xxx.log
```

On `drgXX` nodes, XXX is one of: sda, sdb, sdc, sdd.
On the `dragnet` head node, XXX is one of sg1, sg2 (requires kernel module `sg` loaded, but this is the case by default).

Note that the `dragproc` node has a RAID controller that needs some magic:

```
$ for i in 10 11 9 8 12 13 15 14; do echo "Device ID=$i" && sudo smartctl --
all -d megaraid,$i /dev/sda; done >> smartctl-xxx.log
```

These numbers correspond to the RAID controller's connected drive IDs. In case you want to check or if it changed, you can display these IDs with a special tool called `storcli`; see Alexander's shell log at the end of this wiki page.

To start an extended S.M.A.R.T. check, run (as root or via `sudo`):

```
$ smartctl --test=long /dev/XXX
```

(XXX as well as `dragproc` magic described above apply here too.)
You may continue working, send the same command to the next drive, or even reboot. The drive will continue when no other requests are there. An extended test may take many hours. Afterwards, you can show the S.M.A.R.T. information again to see the result, but if a failure threshold was exceeded, our `smartd` service also sends an e-mail.

To run a sudo command for all drives on all `drgXX` nodes, make a script (ansible -a only takes trivial commands), such as:

```
#!/bin/sh
hostname; smartctl --test=long /dev/sda; smartctl --test=long /dev/sdb;
smartctl --test=long /dev/sdc; smartctl --test=long /dev/sdd
```

mark it executable and run it via ansible as superuser:

```
$ chmod 755 $HOME/yourscript.sh
$ ansible workers -b -K -f 25 -a '$HOME/yourscript.sh'  # + possibly
redirecting stdout and/or stderr
```

then type your password once.

Upon Drive Failure

Tip 1: when drives fail, first:
1. remount the covering partition read-only:

```
$ sudo mount -o remount,ro /dev/mdX  # where mdX is the software RAID
partition that covers your suspicious drive. You can find X using the
''lsblk'' command.
```

The system may protest if files are open in read-write mode. It depends on the situation what to do (ask, kill, wait), but you can find the offending process using the `lsof` command (typically as root; Note: long output!).
2. backup data elsewhere (duh)

Tip 2: when 1 drive fails, run an extended S.M.A.R.T. check on all drives on the cluster to minimize the number of support requests and data center visits.

Important: All drives have a Device Model nr and a Serial Number, etc. Ensure you (Cpt. obvious):

- first report the right IDs + copy-paste the S.M.A.R.T. logs from the right drive
- then replace the right drive in the data center(!) (ask Mike Sipior how to flash the right drive

> lights, etc)
- then return ship the drive(s) with the reported ID(s)!

Then you need to:

- restore the broken RAID (on drgXX) (ask Mike Sipior, or use the `mdadm` command (display of healthy state shown below), then instruct kernel to reload partition tables using the `partprobe` command)
- remount the partition (/dev/mdX to /data[12])
- copy back the backed up data
- once you're sure the backup is no longer needed, delete it

Shipping and Return

If accepted, our vendor generally sends a replacement within 1 working day, provided they have replacements in stock. They also send a UPS label for easy return shipping of the defective component (to their return address in Amsterdam).

Work with Derk Kuipers or someone else at the ASTRON stockroom to send (and pick up) packages. If you have no (valid) UPS label, get a blank label from the reception.

Data Center Visit

For disk replacement, one person is enough. If the node has to be opened up, 2 persons are needed (sliding rails issue + node weight + limited space in data center).

Overview Examples of mdadm and lsblk

How it should look on drgXX nodes, example on drg22 from Aug 2017):

```
[amesfoort@drg22 ~]$ sudo mdadm --detail /dev/md1 /dev/md2
/dev/md1:
          Version : 1.2
    Creation Time : Thu Jul 16 12:13:52 2015
       Raid Level : raid0
       Array Size : 7729802240 (7371.71 GiB 7915.32 GB)
     Raid Devices : 2
    Total Devices : 2
      Persistence : Superblock is persistent

      Update Time : Thu Jul 16 12:13:52 2015
            State : clean
    Active Devices : 2
   Working Devices : 2
    Failed Devices : 0
     Spare Devices : 0
```

```
      Chunk Size : 512K

            Name : drg22.control.lofar:1
            UUID : 071ef430:bfc588c7:8cd42a51:8c53c707
          Events : 0

    Number   Major   Minor   RaidDevice State
       0        8        4        0      active sync   /dev/sda4
       1        8       18        1      active sync   /dev/sdb2
/dev/md2:
         Version : 1.2
   Creation Time : Thu Jul 16 12:13:39 2015
      Raid Level : raid0
      Array Size : 7729802240 (7371.71 GiB 7915.32 GB)
    Raid Devices : 2
   Total Devices : 2
     Persistence : Superblock is persistent

     Update Time : Thu Jul 16 12:13:39 2015
           State : clean
  Active Devices : 2
 Working Devices : 2
  Failed Devices : 0
   Spare Devices : 0

      Chunk Size : 512K

            Name : drg22.control.lofar:2
            UUID : 5690cb47:8d7ff4d4:72cb3108:6857324f
          Events : 0

    Number   Major   Minor   RaidDevice State
       0        8       36        0      active sync   /dev/sdc4
       1        8       50        1      active sync   /dev/sdd2
```

```
[amesfoort@drg22 ~]$ lsblk
NAME      MAJ:MIN RM   SIZE RO TYPE  MOUNTPOINT
sda        8:0     0   3.7T  0 disk
├─sda1     8:1     0     1M  0 part
├─sda2     8:2     0  48.8G  0 part
│ └─md0    9:0     0  48.8G  0 raid1 /
├─sda3     8:3     0  15.6G  0 part  [SWAP]
└─sda4     8:4     0   3.6T  0 part
  └─md1    9:1     0   7.2T  0 raid0 /data1
sdb        8:16    0   3.7T  0 disk
├─sdb1     8:17    0  15.6G  0 part  [SWAP]
└─sdb2     8:18    0   3.6T  0 part
  └─md1    9:1     0   7.2T  0 raid0 /data1
sdc        8:32    0   3.7T  0 disk
├─sdc1     8:33    0     1M  0 part
├─sdc2     8:34    0  48.8G  0 part
```

```
│  └─md0     9:0    0 48.8G  0 raid1 /
├─sdc3      8:35   0 15.6G  0 part  [SWAP]
└─sdc4      8:36   0  3.6T  0 part
  └─md2     9:2    0  7.2T  0 raid0 /data2
sdd         8:48   0  3.7T  0 disk
├─sdd1      8:49   0 15.6G  0 part  [SWAP]
└─sdd2      8:50   0  3.6T  0 part
  └─md2     9:2    0  7.2T  0 raid0 /data2
sr0        11:0    1 1024M  0 rom
```

Using storcli to find RAID controller DIDs

The storcli (and storcli64) utils are from the RAID controller vendor LSI. I used this command to find the Device IDs (DID column) on dragproc for the smartctl command listed earlier on this page.

```
[amesfoort@dragproc storcli]$ pwd
/home/amesfoort/pkg/storcli_all_os/Linux/opt/MegaRAID/storcli
[amesfoort@dragproc storcli]$ sudo ./storcli64 /c0 /eall /sall show
[sudo] password for amesfoort:
Controller = 0
Status = Success
Description = Show Drive Information Succeeded.


Drive Information :
==================


----------------------------------------------------------------------
EID:Slt DID State DG     Size Intf Med SED PI SeSz Model           Sp
----------------------------------------------------------------------
252:0    10 Onln   0 3.637 TB SATA HDD N   N  512B TOSHIBA MC04ACA400E U
252:1    11 Onln   0 3.637 TB SATA HDD N   N  512B TOSHIBA MC04ACA400E U
252:2     9 Onln   0 3.637 TB SATA HDD N   N  512B TOSHIBA MC04ACA400E U
252:3     8 Onln   0 3.637 TB SATA HDD N   N  512B TOSHIBA MC04ACA400E U
252:4    12 Onln   0 3.637 TB SATA HDD N   N  512B TOSHIBA MC04ACA400E U
252:5    13 Onln   0 3.637 TB SATA HDD N   N  512B TOSHIBA MC04ACA400E U
252:6    15 Onln   0 3.637 TB SATA HDD N   N  512B TOSHIBA MC04ACA400E U
252:7    14 Onln   0 3.637 TB SATA HDD N   N  512B TOSHIBA MC04ACA400E U
----------------------------------------------------------------------

EID-Enclosure Device ID|Slt-Slot No.|DID-Device ID|DG-DriveGroup
DHS-Dedicated Hot Spare|UGood-Unconfigured Good|GHS-Global Hotspare
UBad-Unconfigured Bad|Onln-Online|Offln-Offline|Intf-Interface
Med-Media Type|SED-Self Encryptive Drive|PI-Protection Info
SeSz-Sector Size|Sp-Spun|U-Up|D-Down|T-Transition|F-Foreign
UGUnsp-Unsupported|UGShld-UnConfigured shielded|HSPShld-Hotspare shielded
CFShld-Configured shielded|Cpybck-CopyBack|CBShld-Copyback Shielded
```

From:
https://www.astron.nl/lofarwiki/ - **LOFAR *Wiki***

Permanent link:
https://www.astron.nl/lofarwiki/doku.php?id=dragnet:cluster_support&rev=1503913362

Last update: **2017-08-28 09:42**