



Disk Troubleshooting under Linux

Fred Mora

Why troubleshoot disks?

- Disks wear out
- They can sometimes be recovered after partial failure
- Not for the faint of heart
- The best way to avoid trouble:
 - Keep backup on different media
 - Automate the backup process
- You still have to understand when a disk has to be replaced.

Failure modes

- Spinners:
 - Surface wears out
 - Heads scratch the surface
 - Surface defects spread
- SSD
 - Cells wear out
 - Silicon defects appear
- All disks:
 - Controller breaks down
 - Firmware bugs trigger data loss
 - Connectors break (USB disks)

How do I detect a disk failure?

- System unresponsive or very slow
- Abnormal periodic noises from a drive
- Error messages in the log
- Read or write failure in application
- File system corruption, or FS remounted as read-only.

First steps

- After you suspect a disk failure, diag it quickly
- First use `lsscsi` to check the disk controller still responds. For LVM and RAID setups, use `blkid`

```
[fred@desktop] $ lsscsi
[4:0:0:0]    disk      ATA          WDC WD1003FBYX-0 01.0  /dev/sda
[5:0:0:0]    disk      ATA          WDC WD1003FBYX-0 01.0  /dev/sdb
```

```
[fred@server] # blkid
/dev/sdd1: UUID="t00htp-ZJbS-ZTvE-pYdh-vaEC-tZ3s-3cU1Hp"
TYPE="LVM2_member"
/dev/sdc1: UUID="BP8AHh-JcQr-Wt7z-kU6P-K0l3-bkLM-SQj9ik"
TYPE="LVM2_member"
/dev/mapper/vg_docker_1-lv_docker_1: UUID="921b6e78-3604-4578-940d-b2d1942beb16" TYPE="xfs"
```

- Got a response from the disk? Then the controller is still OK. On to the second step!

Media check

- Use the smartctl command to see the disk logs
- The SMART protocol is a bad misnomer
- What you see is a very limited event log and statistics collected during the disk lifetime
- This is NOT an instant disk diagnostic.
- Some SMART attributes are standard
- Others are proprietary and documented only by manufacturers, if at all.

SMART self-test: short test

- Use the SMART command to run self-tests
- All modern disks support the quick and extended tests.
- Short test example:

```
# smartctl -t short /dev/sda
```

```
=== START OF OFFLINE IMMEDIATE AND SELF-TEST SECTION ===  
Sending command: "Execute SMART Short self-test routine  
immediately in off-line mode".  
Drive command "Execute SMART Short self-test routine  
immediately in off-line mode" successful.  
Testing has begun.  
Please wait 2 minutes for test to complete.  
Test will complete after Wed Mar  2 14:55:17 2016
```

SMART self-test: extended test

- Long test is similar, only takes more time and does a surface check:

```
# smartctl -t long /dev/sda
```

```
=== START OF OFFLINE IMMEDIATE AND SELF-TEST SECTION ===
```

```
Sending command: "Execute SMART Extended self-test  
routine immediately in off-line mode".
```

```
Drive command "Execute SMART Extended self-test routine  
immediately in off-line mode" successful.
```

```
Testing has begun.
```

```
Please wait 164 minutes for test to complete.
```

```
Test will complete after Wed Mar  2 17:40:28 2016
```

```
Use smartctl -X to abort test.
```


Example 1: Spinner

- What can SMART tell us about a regular spinning drive?

```
# smartctl -a /dev/sda | less
=== START OF INFORMATION SECTION ===
Model Family:      Western Digital RE4
Device Model:      WDC WD1003FBYX-01Y7B1
Serial Number:     WD-WCAW34838984
LU WWN Device Id:  5 0014ee 2b27f97e8
Firmware Version:  01.01V02
User Capacity:     1,000,204,886,016 bytes [1.00 TB]
Sector Size:       512 bytes logical/physical
Rotation Rate:     7200 rpm
Device is:         In smartctl database [for details use: -P show]
ATA Version is:    ATA8-ACS (minor revision not indicated)
SATA Version is:   SATA 3.0, 3.0 Gb/s (current: 3.0 Gb/s)
Local Time is:     Wed Mar  2 14:25:56 2016 EST
SMART support is:  Available - device has SMART capability.
SMART support is:  Enabled
```

Example 1: Spinner (cont'd)

=== START OF READ SMART DATA SECTION ===

SMART overall-health self-assessment test result: PASSED

Self-test execution status: (0) The previous self-test routine completed without error or no self-test has ever been run.

Total time to complete Offline data collection: (16680) seconds.

Offline data collection capabilities: (0x7b) SMART execute Offline immediate.
Auto Offline data collection on/off support.
Suspend Offline collection upon new command.
Offline surface scan supported.
Self-test supported.
Conveyance Self-test supported.
Selective Self-test supported.

Short self-test routine recommended polling time: (2) minutes.

Extended self-test routine recommended polling time: (164) minutes.

Conveyance self-test routine recommended polling time: (5) minutes.

So far so good. Note the test time values. Extended test takes almost 3 hours.

Example 1: Spinner (cont'd)

Vendor Specific SMART Attributes with Thresholds:

ID#	ATTRIBUTE_NAME	FLAG	VALUE	WORST	THRESH	TYPE	UPDATED	WHEN_FAILED	RAW_VALUE
1	Raw_Read_Error_Rate	0x002f	200	200	051	Pre-fail	Always	-	0
3	Spin_Up_Time	0x0027	174	173	021	Pre-fail	Always	-	4291
4	Start_Stop_Count	0x0032	100	100	000	Old_age	Always	-	65
5	Reallocated_Sector_Ct	0x0033	200	200	140	Pre-fail	Always	-	0
7	Seek_Error_Rate	0x002e	200	200	000	Old_age	Always	-	0
9	Power_On_Hours	0x0032	070	070	000	Old_age	Always	-	21926
10	Spin_Retry_Count	0x0032	100	253	000	Old_age	Always	-	0
11	Calibration_Retry_Count	0x0032	100	253	000	Old_age	Always	-	0
12	Power_Cycle_Count	0x0032	100	100	000	Old_age	Always	-	58
192	Power-Off_Retract_Count	0x0032	200	200	000	Old_age	Always	-	41
193	Load_Cycle_Count	0x0032	200	200	000	Old_age	Always	-	23
194	Temperature_Celsius	0x0022	113	098	000	Old_age	Always	-	34
196	Reallocated_Event_Count	0x0032	200	200	000	Old_age	Always	-	0
197	Current_Pending_Sector	0x0032	200	200	000	Old_age	Always	-	0
198	Offline_Uncorrectable	0x0030	200	200	000	Old_age	Offline	-	0
199	UDMA_CRC_Error_Count	0x0032	200	200	000	Old_age	Always	-	0
200	Multi_Zone_Error_Rate	0x0008	200	200	000	Old_age	Offline	-	0

Here are the dreaded SMART attributes. Those are pretty standard. 196-200 are the money line (as in, a non-zero raw value there means you need to buy a new drive).

Oh, and 21k hours! Good old 1TB Enterprise-grade WD spinner.

Example 1: Spinner (cont'd)

SMART Error Log Version: 1
No Errors Logged

SMART Self-test log structure revision number 1

Num	Test_Description	Status	Remaining	LifeTime(hours)	LBA_of_first_error
# 1	Extended offline	Completed without error	00%	21915	-
# 2	Extended offline	Completed without error	00%	21891	

Here is the test log. We run an automated extended test every day or so. No errors so far. Catches 60% or so of the problems.

Be on the lookout for...

- An extended test showing bad sector
- The SMART log showing ATA command errors
- Attributes that show errors or anomalies
- Example: A 2 TB Passport Green drive that we recently RMA'd

Vendor Specific SMART Attributes with Thresholds:

ID#	ATTRIBUTE_NAME	FLAG	VALUE	WORST	THRESH	TYPE	UPDATED	WHEN_FAILED	RAW_VALUE
1	Raw_Read_Error_Rate	0x002f	200	200	051	Pre-fail	Always	-	0
5	Reallocated_Sector_Ct	0x0033	200	200	140	Pre-fail	Always	-	0
7	Seek_Error_Rate	0x002e	200	200	000	Old_age	Always	-	0
10	Spin_Retry_Count	0x0032	100	253	000	Old_age	Always	-	0
11	Calibration_Retry_Count	0x0032	100	253	000	Old_age	Always	-	0
194	Temperature_Celsius	0x0022	113	098	000	Old_age	Always	-	45
196	Reallocated_Event_Count	0x0032	200	200	000	Old_age	Always	-	0
197	Current_Pending_Sector	0x0032	198	198	000	Old_age	Always	-	948
198	Offline_Uncorrectable	0x0030	198	198	000	Old_age	Offline	-	927
199	UDMA_CRC_Error_Count	0x0032	200	200	000	Old_age	Always	-	0
200	Multi_Zone_Error_Rate	0x0008	200	195	000	Old_age	Offline	-	55

SMART and SSDs

- Since it was not confusing enough, SSDs redefine SMART attributes
- Several attributes in spinners are useless in SSDs
- The failure modes are difference
- Some manufacturers (e.g., Kingston) have good doc on their sites explaining what the attributes are
- Others (SanDisk) actively participate in the smartmontools list and update the DB.
- Example from a Kingston SSD:

195	ECC_Uncorr_Error_Count	0x001c	120	120	000	Old_age	Offline	-	0/189272592
196	Reallocated_Event_Count	0x0033	100	100	003	Pre-fail	Always	-	0
201	Unc_Soft_Read_Err_Rate	0x001c	120	120	000	Old_age	Offline	-	0/189272592
204	Soft_ECC_Correct_Rate	0x001c	120	120	000	Old_age	Offline	-	0/189272592
230	Life_Curve_Status	0x0013	100	100	000	Pre-fail	Always	-	100
231	SSD_Life_Left	0x0013	076	076	010	Pre-fail	Always	-	12884901888
233	SandForce_Internal	0x0032	000	000	000	Old_age	Always	-	98214
234	SandForce_Internal	0x0032	000	000	000	Old_age	Always	-	97162
241	Lifetime_Writes_GiB	0x0032	000	000	000	Old_age	Always	-	97162
242	Lifetime_Reads_GiB	0x0032	000	000	000	Old_age	Always	-	22982



Questions?