



Hard Drive Introspection

Serge Y. Stroobandt

Copyright 2013–2022, licensed under [Creative Commons BY-NC-SA](#)

TODO

- binding text

Introduction



Not so long ago, I came across a badly configured laptop hard drive. On battery power, it was spinning up and down almost every half minute. This and other instances, lent themselves as a great opportunity to discover the vast array of freely available GNU/Linux hard drive tools.

Listing block devices

```
$ lsblk
  NAME MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
  sda   8:0    0 298.1G  0 disk
  └─sda1 8:1    0   512M  0 part /boot
  └─sda2 8:2    0    4.5G  0 part
  └─sda3 8:3    0    44G   0 part /
  └─sda4 8:4    0 249.1G  0 part /home
  sr0   11:0    1  1024M  0 rom

$ lsblk -o NAME,TYPE,SIZE,MODEL,SERIAL,WWN,MOUNTPOINT
  NAME          TYPE  SIZE MODEL          SERIAL
  WWN          MOUNTPOINT
  nvme0n1      disk 931.5G Samsung SSD 970 EVO 1TB S467NX0MB16745H
  eui.0025385b91b14537
  └─nvme0n1p1  part 232.9G
  eui.0025385b91b14537 /
  └─nvme0n1p2  part 698.6G
  eui.0025385b91b14537 /home
```

Listing by ID

```
$ ls -l /dev/disk/by-id
total 0
...
lrwxrwxrwx 1 root root 9 Nov 13 23:04 scsi-SATA_WDC_WD10EARS-00_WD-
WCAV56524564 -> ../../sda
lrwxrwxrwx 1 root root 10 Nov 11 10:48 scsi-SATA_WDC_WD10EARS-00_WD-
WCAV56524564-part1 -> ../../sda1
lrwxrwxrwx 1 root root 10 Nov 11 10:48 scsi-SATA_WDC_WD10EARS-00_WD-
WCAV56524564-part9 -> ../../sda9
...

$ ls -dl --color=always /dev/disk/by-id/* |sed s:^.*\/dev:/dev:
...
/dev/disk/by-id/ata-WDC_WD10EADS-00L5B1_WD-WCAU4D512265 -> ../../sda
/dev/disk/by-id/ata-WDC_WD10EADS-00L5B1_WD-WCAU4D512265-part1 ->
../../sda1
/dev/disk/by-id/ata-WDC_WD10EADS-00L5B1_WD-WCAU4D512265-part9 ->
../../sda9
...
/dev/disk/by-id/wnn-0x50014ee1575de8a2 -> ../../sda
/dev/disk/by-id/wnn-0x50014ee1575de8a2-part1 -> ../../sda1
/dev/disk/by-id/wnn-0x50014ee1575de8a2-part9 -> ../../sda9
```

Note that a drive or a drive partition can have more than one `by-id`. Apart from the ID based on the brand, model name and the serial number, nowadays there might also be a `wnn-` ID. This is the unique **World Wide Name (WWN)** and is also printed on the drive case.

Both type of IDs work fine with ZFS, but the WWN is a bit less telling. If these WWN IDs are *not* referenced by the production system (e.g. a ZFS that has not been exported yet), these may simply be removed with `sudo rm wnn-*`. Trust me; I have done that. Nothing can go wrong as long as the ZFS is in an exported state before doing this. After all, WWN IDs are mere symbolic links to `sd` devices that are created at drive detection. They will automatically reappear when the system is rebooted. Internally, Linux always references `sd` devices.

Listing UUIDs

Universally Unique Identifier (UUID)

```
$ ls -l /dev/disk/by-uuid
total 0
lrwxrwxrwx 1 root root 10 Sep  1 13:08 438b7a73-11c5-466c-
b795-fcf6d0997056 -> ../../sda2
lrwxrwxrwx 1 root root 10 Sep  1 13:08 7c581345-6850-44df-
a387-cdf62cb09bba -> ../../sda4
lrwxrwxrwx 1 root root 10 Sep  1 13:08
8ef04775-ec47-4663-bc77-7e6f8980e195 -> ../../sda3
lrwxrwxrwx 1 root root 10 Sep  1 13:08
9b5c2cd5-43de-4a56-9da9-741710d28dc3 -> ../../sda1
```

```
$ sudo apt install util-linux
$ sudo blkid /dev/sda2
/dev/sda2: UUID="438b7a73-11c5-466c-b795-fcf6d0997056" TYPE="swap"
PARTUUID="00035ac0-02"
```

Physical identification

For the physical identification using storage enclosure LEDs, I created the following bash script:

```
#!/usr/bin/env bash

# https://serverfault.com/a/1108701/175321

if [[ $# -gt 0 ]]
then
    while true
    do
        dd if=$1 of=/dev/null >/dev/null 2>&1 || sudo dd if=$1 of=/dev/null
        >/dev/null 2>&1
        sleep 1
    done
else
    echo -e '\nThis command requires a /dev argument.\n'
fi
```

Unlike `ledctl` from the `ledmon` package, this script also works fine with non-Intel hard drive controllers.

Free disk space

The `df` command is used to query the amount of free disk space. The `-h` option is necessary to print the amount of free disk space in a human-readable fashion. Only local file systems are shown with the `-l` option. To show file system types, use the `-T` option.

```
$ df -hLT
Filesystem      Type      Size  Used Avail Use% Mounted on
udev            devtmpfs  1.9G   0    1.9G   0% /dev
tmpfs           tmpfs     379M   2.5M 376M   1% /run
/dev/sda1       ext4      32G   20G   11G   65% /
tmpfs           tmpfs     1.9G   4.0K 1.9G   1% /dev/shm
tmpfs           tmpfs     5.0M   4.0K 5.0M   1% /run/lock
tmpfs           tmpfs     1.9G   0    1.9G   0% /sys/fs/cgroup
/dev/sda2       ext4      79G   18G   58G   24% /home
tmpfs           tmpfs     379M   12K 379M   1% /run/user/1000
/home/serge/.Private ecryptfs  79G   18G   58G   24% /home/serge
```

Similar, but nicer and in colours:

```
$ dfc -dLT
```

```
$ cat .bash_aliases
alias df='df -hLT'
alias dfc='dfc -dLT'
alias du='du -h'
```

To show the packages consuming the most disk space, use either the `dpigs` command or the `wajig large` command:

```
$ sudo apt install debian-goodies
$ dpigs
```

```
$ sudo apt install wajig
$ wajig large
```

Purge removed packages

Files of previously removed packages may keep on lingering on the system. The following two commands will purge the system of these unnecessary files.

```
$ sudo apt autoremove
$ dpkg --get-selections |grep "^rc" |cut -d " " -f 3 |xargs sudo dpkg --purge
```

SMART monitoring

The acronym **SMART** stands for *Self-Monitoring, Analysis, and Reporting Technology*. Here is how to check whether SMART is supported by a particular drive and whether SMART is switched on.

```
$ sudo apt install smartmontools
$ sudo smartctl -i /dev/sda
    smartctl 6.5 2016-01-24 r4214 [x86_64-linux-4.4.0-34-generic] (local
build)
    Copyright (C) 2002-16, Bruce Allen, Christian Franke,
www.smartmontools.org

    === START OF INFORMATION SECTION ===
    Model Family:      Seagate Barracuda 7200.8
    Device Model:      ST3300831AS
    Serial Number:     4NF0WW20
    Firmware Version:  3.03
    User Capacity:     300,069,052,416 bytes [300 GB]
    Sector Size:       512 bytes logical/physical
    Device is:         In smartctl database [for details use: -P show]
    ATA Version is:    ATA/ATAPI-7 (minor revision not indicated)
    Local Time is:     Wed Aug 24 00:02:09 2016 CEST
    SMART support is: Available - device has SMART capability.
    SMART support is: Enabled
```

Enabling SMART

```
$ sudo smartctl -s on /dev/sda
    smartctl 6.5 2016-01-24 r4214 [x86_64-linux-4.4.0-34-generic] (local
build)
    Copyright (C) 2002-16, Bruce Allen, Christian Franke,
www.smartmontools.org

    === START OF ENABLE/DISABLE COMMANDS SECTION ===
    SMART Enabled.
```

SMART health assessment

```
$ sudo smartctl -H /dev/sda
    smartctl 6.5 2016-01-24 r4214 [x86_64-linux-4.4.0-34-generic] (local
build)
    Copyright (C) 2002-16, Bruce Allen, Christian Franke,
www.smartmontools.org

    === START OF READ SMART DATA SECTION ===
    SMART overall-health self-assessment test result: PASSED
```

All SMART output

```
$ sudo smartctl -a /dev/sda
smartctl 6.5 2016-01-24 r4214 [x86_64-linux-4.4.0-34-generic] (local
build)
Copyright (C) 2002-16, Bruce Allen, Christian Franke,
www.smartmontools.org

=== START OF INFORMATION SECTION ===
Model Family:      Seagate Barracuda 7200.8
Device Model:      ST3300831AS
Serial Number:     4NF0WW20
Firmware Version: 3.03
User Capacity:     300,069,052,416 bytes [300 GB]
Sector Size:       512 bytes logical/physical
Device is:         In smartctl database [for details use: -P show]
ATA Version is:    ATA/ATAPI-7 (minor revision not indicated)
Local Time is:     Wed Aug 24 00:02:17 2016 CEST
SMART support is: Available - device has SMART capability.
SMART support is: Enabled

=== START OF READ SMART DATA SECTION ===
SMART overall-health self-assessment test result: PASSED

General SMART Values:
Offline data collection status: (0x82) Offline data collection activity
was completed without error.
Auto Offline Data Collection: Enabled.
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:      ( 430) seconds.
Offline data collection
capabilities:          (0x5b) SMART execute Offline immediate.
Auto Offline data collection on/off support.
Suspend Offline collection upon new
command.
Offline surface scan supported.
Self-test supported.
No Conveyance Self-test supported.
Selective Self-test supported.
SMART capabilities:    (0x0003) Saves SMART data before entering
power-saving mode.
Supports SMART auto save timer.
Error logging capability: (0x01) Error logging supported.
General Purpose Logging supported.

Short self-test routine
recommended polling time:      (   1) minutes.
Extended self-test routine
recommended polling time:      ( 101) minutes.
```

```

SMART Attributes Data Structure revision number: 10
Vendor Specific SMART Attributes with Thresholds:
ID# ATTRIBUTE_NAME          FLAG         VALUE WORST THRESH TYPE
UPDATED  WHEN_FAILED RAW_VALUE
  1 Raw_Read_Error_Rate     0x000f       063   057   006   Pre-fail
Always    -          162822491
  3 Spin_Up_Time            0x0003       097   097   000   Pre-fail
Always    -              0
  4 Start_Stop_Count        0x0032       099   099   020   Old_age
Always    -          1913
  5 Reallocated_Sector_Ct   0x0033       100   100   036   Pre-fail
Always    -              0
  7 Seek_Error_Rate         0x000f       084   060   030   Pre-fail
Always    -      269974882
  9 Power_On_Hours          0x0032       090   090   000   Old_age
Always    -          8792
 10 Spin_Retry_Count        0x0013       100   100   097   Pre-fail
Always    -              0
 12 Power_Cycle_Count       0x0032       099   099   020   Old_age
Always    -          1920
194 Temperature_Celsius    0x0022       044   052   000   Old_age
Always    -    44 (0 14 0 0 0)
 195 Hardware_ECC_Recovered 0x001a       063   057   000   Old_age
Always    -      162822491
 197 Current_Pending_Sector 0x0012       100   100   000   Old_age
Always    -              0
 198 Offline_Uncorrectable   0x0010       100   100   000   Old_age
Offline   -              0
 199 UDMA_CRC_Error_Count   0x003e       200   200   000   Old_age
Always    -              0
 200 Multi_Zone_Error_Rate   0x0000       100   253   000   Old_age
Offline   -              0
 202 Data_Address_Mark_Errs 0x0032       100   253   000   Old_age
Always    -              0

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 Short offline          Completed without error     00%
0 -

SMART Selective self-test log data structure revision number 1
SPAN  MIN_LBA  MAX_LBA  CURRENT_TEST_STATUS
  1      0        0  Not_testing
  2      0        0  Not_testing
  3      0        0  Not_testing
  4      0        0  Not_testing
  5      0        0  Not_testing

Selective self-test flags (0x0):
  After scanning selected spans, do NOT read-scan remainder of disk.
  If Selective self-test is pending on power-up, resume after 0 minute
  delay.

```

hddtemp

```
$ sudo apt install hddtemp

$ sudo hddtemp /dev/sd[abcd]
/dev/sda: WDC WD10EARS-00Y5B1: 33°C
/dev/sdb: WDC WD10EADS-00L5B1: 31°C
/dev/sdc: WDC WD10EADS-00M2B0: 33°C
/dev/sdd: Generic-SD/MMC: S.M.A.R.T. not available

$ sudo hddtemp /dev/sd*[^[:digit:]] 2>/dev/null |sort -k3rn
/dev/sda: Generic MassStorageClass: S.M.A.R.T. not available
/dev/sde: WDC WD20EFRX-68EUZN0: 35°C
/dev/sdf: WDC WD20EFRX-68EUZN0: 34°C
/dev/sdg: WDC WD20EFRX-68EUZN0: 36°C
```

The latter `hddtemp` evocation is the most useful and may be added as the default to `.bash_aliases` as follows:

```
alias hddtemp='sudo hddtemp /dev/sd*[^[:digit:]] 2>/dev/null |sort -k3rn'
```

hdparm

- [Wikipedia.org](https://en.wikipedia.org/wiki/Hdparm)
- [hdparm: Out of respect for times past](#)

Installing hdparm

On modern Debian derived distributions, `hdparm` can be installed as follows:

```
$ sudo apt install hdparm
```

Detailed drive info from hdparm

[man page](#)

`-I`: Request identification info directly from the drive, which is displayed in a new expanded format with considerably more detail than with the older `-i` flag.

```
$ sudo hdparm -I /dev/sda
```


/dev/sda:

ATA device, with non-removable media

Model Number: WDC WD2500BEVT-08A23T1

Serial Number: WD-WXK1A8004345

Firmware Revision: 02.01A02

Transport: Serial, SATA 1.0a, SATA II Extensions, SATA Rev 2.5, SATA Rev 2.6

Standards:

Supported: 8 7 6 5

Likely used: 8

Configuration:

Logical max current

cylinders 16383 16383

heads 16 16

sectors/track 63 63

--

CHS current addressable sectors: 16514064

LBA user addressable sectors: 268435455

LBA48 user addressable sectors: 488397168

Logical/Physical Sector size: 512 bytes

device size with M = 1024*1024: 238475 MBytes

device size with M = 1000*1000: 250059 MBytes (250 GB)

cache/buffer size = 8192 KBytes

Nominal Media Rotation Rate: 5400

Capabilities:

LBA, IORDY(can be disabled)

Queue depth: 32

Standby timer values: spec'd by Standard, no device specific minimum

R/W multiple sector transfer: Max = 16 Current = 16

Advanced power management level: 127

Recommended acoustic management value: 128, current value: 128

DMA: mdma0 mdma1 mdma2 udma0 udma1 udma2 udma3 udma4 udma5 *udma6

Cycle time: min=120ns recommended=120ns

PIO: pio0 pio1 pio2 pio3 pio4

Cycle time: no flow control=120ns IORDY flow control=120ns

Commands/features:

Enabled Supported:

- * SMART feature set
- Security Mode feature set
- * Power Management feature set
- * Write cache
- * Look-ahead
- * Host Protected Area feature set
- * WRITE_BUFFER command
- * READ_BUFFER command
- * DOWNLOAD_MICROCODE
- * Advanced Power Management feature set
- SET_MAX security extension
- * Automatic Acoustic Management feature set
- * 48-bit Address feature set
- * Device Configuration Overlay feature set
- * Mandatory FLUSH_CACHE
- * FLUSH_CACHE_EXT
- * SMART error logging
- * SMART self-test
- * General Purpose Logging feature set
- * WRITE_{DMA|MULTIPLE}_FUA_EXT
- * 64-bit World wide name
- * IDLE_IMMEDIATE with UNLOAD
- * Disable Data Transfer After Error Detection
- * WRITE_UNCORRECTABLE_EXT command
- * {READ,WRITE}_DMA_EXT_GPL commands
- * Segmented DOWNLOAD_MICROCODE
- * Gen1 signaling speed (1.5Gb/s)
- * Gen2 signaling speed (3.0Gb/s)
- * Native Command Queueing (NCQ)
- * Host-initiated interface power management
- * Phy event counters
- * Idle-Unload when NCQ is active
- * NCQ priority information
- * DMA Setup Auto-Activate optimization
- * Device-initiated interface power management
- * Software settings preservation
- * SMART Command Transport (SCT) feature set
- * SCT Read/Write Long (AC1), obsolete
- * SCT Write Same (AC2)
- * SCT Features Control (AC4)
- * SCT Data Tables (AC5)
- unknown 206[12] (vendor specific)
- unknown 206[13] (vendor specific)
- unknown 206[14] (vendor specific)

Security:

Master password revision code = 65534
supported
not enabled
not locked
frozen
not expired: security count
supported: enhanced erase
64min for SECURITY ERASE UNIT. 64min for ENHANCED SECURITY ERASE

UNIT.

Logical Unit WWN Device Identifier: 50014ee25a524517
NAA : 5
IEEE OUI : 0014ee
Unique ID : 25a524517
Checksum: correct

Querying spin-down parameters with hdparm

Caveat: Refrain from using laptop-mode-tools

Never install `laptop-mode-tools`, regardless whether `hdparm` is used or not! The package `laptop-mode-tools` will overwrite the `hdparm -B` option and set it to 1. Such a setting will cause any hard drive to fail early.

man page

`-B`: Query/set *Advanced Power Management* feature, if the drive supports it. A low value means aggressive power management and a high value means better performance. Values which permit spin-down range from 1 to 127; values 128 through 254 will not permit spin-down. The highest degree of power management is attained with a setting of 1, and the highest I/O performance with a setting of 254. A value of 255 tells `hdparm` to disable *Advanced Power Management* altogether on the drive. Not all drives support disabling it, but most do.

`-S`: Put the drive into idle (low-power) mode, and also set the standby (spin-down) timeout for the drive. This timeout value is used by the drive to determine how long to wait after the last disk activity before turning off the spindle motor to save power. Once spun down, a drive may take as long as 30 seconds to respond to a subsequent disk access, though most drives are much quicker. The encoding of the timeout value is somewhat peculiar:

- The value of 0 disables any timeout; The device will refrain from automatically entering standby mode.
- Values from 1 to 240 specify multiples of 5 seconds, yielding timeouts from 5 seconds to 20 minutes.
- Values from 241 to 251 specify from 1 to 11 units of 30 minutes, yielding timeouts from 30 minutes to 5.5 hours.
- A value of 252 signifies a timeout of 21 minutes.
- A value of 253 sets a vendor-defined timeout period between 8 and 12 hours.
- The value of 254 is reserved.
- A value of 255 is interpreted as 21 minutes plus 15 seconds.

Note that some older drives may have very different interpretations of these values.

Setting spin-down parameters with hdparm

```
$ sudo hdparm -B127 /dev/sda
/dev/sda:
  setting Advanced Power Management level to 0x7f (127)
  APM_level = 127
```

```
$ sudo hdparm -S60 /dev/sda
/dev/sda:
  setting standby to 60 (5 minutes)
```



This work is licensed under a [Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License](https://creativecommons.org/licenses/by-nc-sa/4.0/).

Other licensing available on request.

Unattended [CSS](https://www.cssprint.com/) typesetting with . The Prince logo features the word "Prince" in a stylized purple font with "Print with CSS" written in smaller text above it.

This work is published at <https://hamwaves.com/hd/en/>.

Last update: Tuesday, August 23, 2022.