

D-Link DNS-323 network attached storage

t=0x0015d40 [0,0]

Contents: [Dobrica PavlinuÅ;jiÄ 's random unstructured stuff]

- [Dobrica PavlinuÅ;jiÄ 's random unstructured stuff \(Overview\)](#)
- [Dobrica PavlinuÅ;jiÄ 's random unstructured stuff \(RAID\)](#)
- [Dobrica PavlinuÅ;jiÄ 's random unstructured stuff \(Debian\)](#)
 - ◆ [Dobrica PavlinuÅ;jiÄ 's random unstructured stuff \(2.6.26-2-orion5x\)](#)
 - ◇ [Dobrica PavlinuÅ;jiÄ 's random unstructured stuff \(raid0\)](#)
 - ◆ [Dobrica PavlinuÅ;jiÄ 's random unstructured stuff \(2.6.30\)](#)
 - ◆ [Dobrica PavlinuÅ;jiÄ 's random unstructured stuff \(updates\)](#)

Overview

dmesg.txt

```
dlink-DFDADE:~# uname -a
Linux dlink-DFDADE 2.6.12.6-arm1 #30 Mon Aug 18 14:19:14 CST 2008 armv5tej1 GNU/Linux
```

```
dlink-DFDADE:~# free
              total        used         free       shared    buffers     cached
Mem:           61904        54808         7096           0        11824        30592
-/+ buffers/cache:      12392        49512
Swap:        1060208           0       1060208
```

```
dlink-DFDADE:~# cat /proc/cpuinfo
Processor       : ARM926EJ-Sid(wb) rev 0 (v5l)
BogoMIPS        : 331.77
Features        : swp half thumb fastmult edsp java
CPU implementer : 0x41
CPU architecture: 5TEJ
CPU variant     : 0x0
CPU part        : 0x926
CPU revision    : 0
Cache type      : write-back
Cache clean     : cp15 c7 ops
Cache lockdown  : format C
Cache format    : Harvard
I size          : 32768
I assoc         : 1
I line length   : 32
I sets          : 1024
D size          : 32768
D assoc         : 1
D line length   : 32
D sets          : 1024

Hardware        : MV-88fxx81
Revision        : 0000
Serial          : 0000000000000000
```

RAID

```
Personalities : [linear] [raid0] [raid1]
md0 : active raid0 sda2[0] sdb2[1]
      2925532672 blocks 64k chunks

unused devices: <none>
dlink-DFDADE:~# hdparm -tT /dev/sda /dev/sdb /dev/md0

/dev/sda:
Timing cached reads:   210 MB in  2.01 seconds = 104.48 MB/sec
Timing buffered disk reads:  72 MB in  3.00 seconds = 24.00 MB/sec

/dev/sdb:
Timing cached reads:   212 MB in  2.01 seconds = 105.47 MB/sec
Timing buffered disk reads: 104 MB in  3.03 seconds = 34.32 MB/sec

/dev/md0:
Timing cached reads:   208 MB in  2.01 seconds = 103.48 MB/sec
BLKGETSIZE failed: File too large

dlink-DFDADE:~# dd_rescue /dev/sda /dev/null
dd_rescue: (info): ipos:   1044480.0k, opos:   1044480.0k, xferd:   1044480.0k
                  errs:     0, errxfer:     0.0k, succxfer:  1044480.0k
                  +curr.rate:  37169kB/s, avg.rate:   32123kB/s, avg.load: 60.7%

dlink-DFDADE:~# dd_rescue /dev/md0 /dev/null
dd_rescue: (info): ipos:   729536.0k, opos:   729536.0k, xferd:   729536.0k
                  errs:     0, errxfer:     0.0k, succxfer:  729536.0k
                  +curr.rate:  37176kB/s, avg.rate:   37502kB/s, avg.load: 76.5%
```

Debian

- <http://www.cyrius.com/debian/orion/d-link/dns-323/install.html>

2.6.26-2-orion5x

raid0

```
Linux dlink-DFDADE 2.6.26-2-orion5x #1 Thu Aug 20 05:04:03 UTC 2009 armv5tel GNU/Linux
Personalities : [raid1] [raid6] [raid5] [raid4] [raid0]
md1 : active raid0 dm-1[1] dm-0[0]
      2097024 blocks 64k chunks

md0 : active raid1 sda4[0] sdb4[1]
      513984 blocks [2/2] [UU]

unused devices: <none>

/dev/sda:
Timing cached reads:   212 MB in  2.01 seconds = 105.39 MB/sec
Timing buffered disk reads:  72 MB in  3.05 seconds = 23.58 MB/sec

/dev/sdb:
Timing cached reads:   218 MB in  2.01 seconds = 108.26 MB/sec
Timing buffered disk reads:  88 MB in  3.00 seconds = 29.33 MB/sec

/dev/md1:
```

```
Timing cached reads: 208 MB in 2.01 seconds = 103.46 MB/sec
Timing buffered disk reads: 60 MB in 3.07 seconds = 19.57 MB/sec
```

```
Linux dlink-DFDADE 2.6.26-2-orion5x #1 Thu Aug 20 05:04:03 UTC 2009 armv5tel GNU/Linux
Personalities : [raid1] [raid6] [raid5] [raid4] [raid0]
md1 : active (auto-read-only) raid1 dm-1[1] dm-0[0]
      1048512 blocks [2/2] [UU]
      resync=PENDING
```

```
md0 : active raid1 sda4[0] sdb4[1]
      513984 blocks [2/2] [UU]
```

```
unused devices: <none>
```

```
/dev/sda:
Timing cached reads: 220 MB in 2.00 seconds = 109.73 MB/sec
Timing buffered disk reads: 82 MB in 3.01 seconds = 27.22 MB/sec
```

```
/dev/sdb:
Timing cached reads: 216 MB in 2.02 seconds = 107.12 MB/sec
Timing buffered disk reads: 90 MB in 3.09 seconds = 29.14 MB/sec
```

```
/dev/md1:
Timing cached reads: 208 MB in 2.02 seconds = 103.19 MB/sec
Timing buffered disk reads: 52 MB in 3.05 seconds = 17.06 MB/sec
```

2.6.30

<http://www.cyrius.com/journal/debian/orion/d-link/dns-323/dns-323-fan-control>

```
dlink-DFDADE:~/mdadm# ./test.sh
+ ./remove.md1.sh
+ mdadm --manage --stop /dev/md1
mdadm: error opening /dev/md1: No such file or directory
+ mdadm --manage --remove /dev/md1
mdadm: error opening /dev/md1: No such file or directory
+ yes
+ mdadm --create --verbose /dev/md1 --level=0 --raid-devices=2 --force /dev/vga/raid.a /dev/vgb/raid.b
mdadm: chunk size defaults to 64K
mdadm: /dev/vga/raid.a appears to be part of a raid array:
      level=raid1 devices=2 ctime=Sat Sep 5 14:54:52 2009
mdadm: /dev/vgb/raid.b appears to be part of a raid array:
      level=raid1 devices=2 ctime=Sat Sep 5 14:54:52 2009
Continue creating array? mdadm: array /dev/md1 started.
+ ./hdparm-test.sh
+ test -d out
+ uname -a
+ cat /proc/mdstat
+ hdparm -tT /dev/sda /dev/sdb /dev/md1
++ date +%Y%m%d_%H%M%S
+ tee out/20090905_151032
Linux dlink-DFDADE 2.6.30-1-orion5x #1 Tue Aug 18 04:19:30 UTC 2009 armv5tel GNU/Linux
Personalities : [raid1] [raid0]
md1 : active raid0 dm-0[1] dm-1[0]
      2097024 blocks 64k chunks

md0 : active raid1 sda4[0] sdb4[1]
      513984 blocks [2/2] [UU]
```

unused devices: <none>

/dev/sda:

Timing cached reads: 246 MB in 2.01 seconds = 122.14 MB/sec
Timing buffered disk reads: 132 MB in 3.00 seconds = 43.93 MB/sec

/dev/sdb:

Timing cached reads: 242 MB in 2.01 seconds = 120.27 MB/sec
Timing buffered disk reads: 138 MB in 3.01 seconds = 45.87 MB/sec

/dev/md1:

Timing cached reads: 234 MB in 2.01 seconds = 116.15 MB/sec
Timing buffered disk reads: 130 MB in 3.03 seconds = 42.85 MB/sec

```
+ ./remove.md1.sh
+ mdadm --manage --stop /dev/md1
mdadm: stopped /dev/md1
+ mdadm --manage --remove /dev/md1
+ yes
+ mdadm --create --verbose /dev/md1 --level=1 --raid-devices=2 --spare-devices=0 --force /dev/vga
mdadm: /dev/vga/raid.a appears to be part of a raid array:
    level=raid0 devices=2 ctime=Sat Sep  5 15:10:31 2009
mdadm: /dev/vga/raid.b appears to be part of a raid array:
    level=raid0 devices=2 ctime=Sat Sep  5 15:10:31 2009
mdadm: size set to 1048512K
Continue creating array? mdadm: array /dev/md1 started.
+ ./hdparm-test.sh
+ test -d out
+ uname -a
++ date +%Y%m%d_%H%M%S
+ cat /proc/mdstat
+ hdparm -tT /dev/sda /dev/sdb /dev/md1
+ tee out/20090905_151114
Linux dlink-DFDADE 2.6.30-1-orion5x #1 Tue Aug 18 04:19:30 UTC 2009 armv5tel GNU/Linux
Personalities : [raid1] [raid0]
md1 : active (auto-read-only) raid1 dm-0[1] dm-1[0]
    1048512 blocks [2/2] [UU]
    resync=PENDING

md0 : active raid1 sda4[0] sdb4[1]
    513984 blocks [2/2] [UU]
```

unused devices: <none>

/dev/sda:

Timing cached reads: 116 MB in 2.01 seconds = 57.84 MB/sec
Timing buffered disk reads: 118 MB in 3.00 seconds = 39.32 MB/sec

/dev/sdb:

Timing cached reads: 114 MB in 2.01 seconds = 56.79 MB/sec
Timing buffered disk reads: 140 MB in 3.03 seconds = 46.27 MB/sec

/dev/md1:

Timing cached reads: 234 MB in 2.02 seconds = 116.07 MB/sec
Timing buffered disk reads: 110 MB in 3.01 seconds = 36.49 MB/sec

updates

Martin Michlmayr

- [Research on FOSS foundations](#)

I worked on research on FOSS foundations and published two reports:

Growing Open Source Projects with a Stable Foundation

This primer covers non-technical aspects that the majority of projects will have to consider at some point. It also explains how FOSS foundations can help projects grow and succeed.

This primer explains:

- ◆ What issues and areas to consider
- ◆ How other projects and foundations have approached these topics
- ◆ What FOSS foundations bring to the table
- ◆ How to choose a FOSS foundation

You can download [Growing Open Source Projects with a Stable Foundation](#).

Research report

The research report describes the findings of the research and aims to help understand the operations and challenges FOSS foundations face.

This report covers topics such as:

- ◆ Role and activities of foundations
- ◆ Challenges faced and gaps in the service offerings
- ◆ Operational aspects, including reasons for starting an org and choice of jurisdiction
- ◆ Trends, such as the "foundation in a foundation" model
- ◆ Recommendations for different stakeholders

You can download the [research report](#).

Acknowledgments

This research was sponsored by Ford Foundation and Alfred P. Sloan Foundation. The research was part of their [Critical Digital Infrastructure Research](#) initiative, which investigates the role of open source in digital infrastructure.

- [ledger2beancount 2.6 released](#)

I released version 2.6 of [ledger2beancount](#), a ledger to beancount converter.

Here are the changes in 2.6:

- ◆ Round calculated total if needed for `price==cost` comparison
- ◆ Add `narration_tag` config variable to set narration from metadata
- ◆ Retain unconsummated payee/payer metadata
- ◆ Ensure UTF-8 output and assume UTF-8 input
- ◆ Document UTF-8 issue on Windows systems
- ◆ Add option to move posting-level tags to the transaction itself
- ◆ Add support for the `alias` sub-directive of `account` declarations
- ◆ Add support for the `payee` sub-directive of `account` declarations
- ◆ Support configuration file called `.ledger2beancount.yaml`

- ◆ Fix uninitialised value warning in hledger mode
- ◆ Print warning if account in assertion has sub-accounts
- ◆ Set commodity for commodity-less balance assertion
- ◆ Expand path name of `beancount_header` config variable
- ◆ Document handling of buckets
- ◆ Document pre- and post-processing examples
- ◆ Add `Dockerfile` to create Docker image

Thanks to Alexander Baier, Daniele Nicolodi, and GitHub users bratekarate, faaaf0 and mefromthepast for various bug reports and other input.

Thanks to Dennis Lee for adding a Dockerfile and to Vinod Kurup for fixing a bug.

Thanks to Stefano Zacchiroli for testing.

You can get [ledger2beancount from GitHub](#).

- [beancount2ledger 1.3 released](#)

I released version 1.3 of [beancount2ledger](#), the beancount to ledger converter that was moved from `bean-report ledger` into a standalone tool.

You can get [beancount2ledger from GitHub](#) or via `pip install`.

Here are the changes in 1.3:

- ◆ Add rounding postings only when required ([issue #9](#))
- ◆ Avoid printing too much precision for a currency ([issue #21](#))
- ◆ Avoid creating two or more postings with null amount ([issue #23](#))
- ◆ Add price to cost when needed by ledger ([issue #22](#))
- ◆ Preserve posting order ([issue #18](#))
- ◆ Add config option `indent`
- ◆ Show metadata with hledger output
- ◆ Support setting auxiliary dates and posting dates from metadata ([issue #14](#))
- ◆ Support setting the code of transactions from metadata
- ◆ Support mapping of account and currency names ([issue #24](#))
- ◆ Improve documentation:
 - ◇ Add user guide
 - ◇ Document limitations ([issue #12](#))

- [ledger2beancount 2.5 released](#)

I released version 2.5 of [ledger2beancount](#), a ledger to beancount converter.

Here are the changes in 2.5:

- ◆ Don't create negative cost for lot without cost
- ◆ Support complex implicit conversions
- ◆ Handle typed metadata with value `0` correctly
- ◆ Set per-unit instead of total cost when cost is missing from lot
- ◆ Support commodity-less amounts
- ◆ Convert transactions with no amounts or only `0` amounts to notes
- ◆ Fix parsing of transaction notes
- ◆ Keep tags in transaction notes on same line as transaction header
- ◆ Add beancount config options for non-standard root names automatically
- ◆ Fix conversion of fixated prices to costs

- ◆ Fix removal of price when `price==cost` but when they use different number formats
- ◆ Fix removal of price when `price==cost` but per-unit and total notation mixed
- ◆ Fix detection of tags and metadata after posting/aux date
- ◆ Use `D` directive to set default commodity for `hledger`
- ◆ Improve support for postings with commodity-less amounts
- ◆ Allow empty comments
- ◆ Preserve leading whitespace in comments in postings and transaction headers
- ◆ Preserve indentation for tags and metadata
- ◆ Preserve whitespace between amount and comment
- ◆ Refactor code to use more data structures
- ◆ Remove dependency on `Config::Onion` module

Thanks to input from Remco RÃ³nders, Yuri Khan, and Thierry. Thanks to Stefano Zacchiroli and Kirill Goncharov for testing my changes.

You can get [ledger2beancount from GitHub](#)

- [ledger2beancount 2.4 released](#)

I released version 2.4 of [ledger2beancount](#), a ledger to beancount converter.

There are two notable changes in this release:

1. I fixed two regressions introduced in the last release. Sorry about the breakage!
2. I improved support for `hledger`. I believe all syntax differences in `hledger` are supported now.

Here are the changes in 2.4:

- ◆ Fix regressions introduced in version 2.3
 - ◇ Handle price directives with comments
 - ◇ Don't assume implicit conversion when price is on second posting
- ◆ Improve support for `hledger`
 - ◇ Fix parsing of `hledger` tags
 - ◇ Support commas as decimal markers
 - ◇ Support digit group marks through `commodity` and `D` directives
 - ◇ Support `end aliases` directive
 - ◇ Support regex aliases
 - ◇ Recognise total balance assertions
 - ◇ Recognise sub-account balance assertions
- ◆ Add support for `define` directive
- ◆ Convert all uppercase metadata tags to all lowercase
- ◆ Improve handling of ledger lots without cost
- ◆ Allow transactions without postings
- ◆ Fix parsing issue in commodity declarations
- ◆ Support commodities that contain quotation marks
- ◆ Add `--version` option to show version
- ◆ Document problem of mixing `apply` and `include`

Thanks to Kirill Goncharov for pointing out one regressions, to Taylor R Campbell for for a patch, to Stefano Zacchiroli for some input, and finally to Simon Michael for input on `hledger`!

You can get [ledger2beancount from GitHub](#)

- [beancount2ledger 1.1 released](#)

Martin Blais recently [announced that he'd like to re-organize the beancount code](#) and split out some functionality into separate projects, including the beancount to ledger/hledger

conversion code previously provided by bean-report.

I agreed to take on the maintenance of this code and I've now released [beancount2ledger](#), a [beancount to ledger/hledger converter](#).

You can install beancount2ledger with pip:

```
pip3 install beancount2ledger
```

Please [report issues to the GitHub tracker](#).

There are a number of outstanding issues I'll fix soon, but please report any other issues you encounter.

Note that I'm not very familiar with hledger. I intend to sync up with hledger author Simon Michael soon, but please file an issue if you notice any problems with the hledger conversion.

Version 1.1 contains a number of fixes compared to the latest code in bean-report:

1.1 (2020-07-24)

- ◆ Preserve metadata information ([issue #3](#))
- ◆ Preserve cost information (lot dates and lot labels/notes) ([issue #5](#))
- ◆ Avoid adding two prices in hledger ([issue #2](#))
- ◆ Avoid trailing whitespace in account open declarations ([issue #6](#))
- ◆ Fix indentation issue in postings ([issue #8](#))
- ◆ Fix indentation issue in price entries
- ◆ Drop time information from price (P) entries
- ◆ Add documentation
- ◆ Relicense under GPL-2.0-or-later ([issue #1](#))

1.0 (2020-07-22)

- ◆ Split ledger and hledger conversion from `bean-report` into a standalone tool
- ◆ Add man page for `beancount2ledger(1)`
- [ledger2beancount 2.3 released](#)

I released version 2.3 of [ledger2beancount](#), a ledger to beancount converter.

There are three notable changes with this release:

1. Performance has significantly improved. One large, real-world test case has gone from around 160 seconds to 33 seconds. A smaller test case has gone from 11 seconds to ~3.5 seconds.
2. The documentation is [available online now](#) (via Read the Docs).
3. The [repository has moved](#) to the beancount GitHub organization.

Here are the changes in 2.3:

- ◆ Improve speed of ledger2beancount significantly
- ◆ Improve parsing of postings for accuracy and speed
- ◆ Improve support for inline math
- ◆ Handle lots without cost
- ◆ Fix parsing of lot notes followed by a virtual price

- ◆ Add support for lot value expressions
- ◆ Make parsing of numbers more strict
- ◆ Fix behaviour of dates without year
- ◆ Accept default ledger date formats without configuration
- ◆ Fix implicit conversions with negative prices
- ◆ Convert implicit conversions in a more idiomatic way
- ◆ Avoid introducing trailing whitespace with hledger input
- ◆ Fix loading of config file
- ◆ Skip ledger directive `import`
- ◆ Convert documentation to `mkdocs`

Thanks to Colin Dean for some feedback. Thanks to Stefano Zacchiroli for prompting me into investigating performance issues (and thanks to the developers of the Devel::NYTProf profiler).

You can get [ledger2beancount from GitHub](#)

- [ledger2beancount 2.2 released](#)

I released version 2.2 of [ledger2beancount](#), a ledger to beancount converter.

Here are the changes in 2.2:

- ◆ Show warning for unknown `apply` directive
- ◆ Recognize `apply rate` directive (an alias of `apply fixed`)
- ◆ Don't convert meta-data on ignored virtual postings but keep as comments
- ◆ Update location of beancount repository

You can get [ledger2beancount from GitHub](#).

Thanks to GitHub user MarinBernard for reporting a bug with virtual postings!

- [ledger2beancount 2.1 released](#)

I released version 2.1 of [ledger2beancount](#), a ledger to beancount converter.

Here are the changes in 2.1:

- ◆ Handle postings with posting dates and comments but no amount
- ◆ Show transactions with only one posting (without `bucket`)
- ◆ Adding spacing between automatic declarations
- ◆ Preserve preliminary info at the top

You can get [ledger2beancount from GitHub](#).

Thanks to Thierry (thdox) for reporting a bug and for fixing some typos in the documentation. Thanks to Stefano Zacchiroli for some good feedback.

- [ledger2beancount 2.0 released](#)

I released version 2.0 of [ledger2beancount](#), a ledger to beancount converter.

Here are the changes in 2.0:

- ◆ Handle comments in `account` and `commodity` declarations
- ◆ Handle transactions with a single posting (without `bucket`)
- ◆ Handle empty metadata values
- ◆ Rewrite Emacs modeline

You can get [ledger2beancount from GitHub](#).