

parting ways with partman?

nick black for debconf 2021-08-24

what's a partman?

- d-i component responsible for partitioning disks, creating filesystems/swap/crypto, mounting `/target`, and preparing `/etc/fstab`
 - Does **not** install bootloader
- 20+ udebs, each a collection of shell scripts and descriptor files, small bit of C (~1.5MB on-iso)

[!!!] Partition disks

This is an overview of your currently configured partitions and mount points. Select a partition to modify its settings (file system, mount point, etc.), a free space to create partitions, or a device to initialize its partition table.

Guided partitioning
Configure software RAID
Configure the Logical Volume Manager
Configure encrypted volumes
Configure iSCSI volumes

Virtual disk 1 (vda) - 21.5 GB Virtio Block Device

	1.0 MB			FREE SPACE	
#1	536.9 MB	B	f	ESP	
#2	19.9 GB		f	ext4	/
#3	1.0 GB		f	swap	swap
	1.0 MB			FREE SPACE	

Undo changes to partitions

Finish partitioning and write changes to disk

<Go Back>

- d-i also ships the `parted` udeb, both for manual setup and as a dep of partman
- partman git has history going back to 2004
- guided setup + preseeded headless setup
- well-translated
- declarative, UI-free
 - ties in unobtrusively with the rest of d-i, *but*
 - limits flexibility and power

why replace partman?

- There's no compelling reason to replace partman, and anything hoping to requires quite a bit of functionality, some of it only germane to the system installation use case.
- There would be a significant amount of work building and verifying any replacement, and it's unlikely that anyone would want to do so.
- It's not like a replacement is going to just fall, fully-formed, from the heavens...

...let's take it back to 2012

- Experimental Debian derivative "SprezzOS"
- One goal was well-integrated deployment of ZFS
 - Installer support was mandatory
- Another was proper alignment on SSDs/AF disks
 - This likewise needed install-time support
- I wrote a disk manager and integrated it with d-i
- Admitted to Debian archive in 2020



- Doesn't rely on d-i infrastructure, and is thus also available as a post-install utility...
- ...but explicitly designed to fit into the partman-sized hole in d-i.
- `--target` option runs in *install mode* where successful return implies working `/etc/fstab`, ESP/MBR, installed bootloader.
- libreadline and TUI modes (readline-based is probably best for Braille/screenreaders).

```

Growlight Blockdevs Partitions Info
-----[ahci-0 (256Gbps to chip, 12Gbps (4%) demanded)]-----[-]
  sdm
  * 7200 rpm me1111111111111111 zfs_member "chungus" (12.00T) 111111111111119em
  29° no i/o | ST12000NM0007-2A SN03 12.00T 4096B gpt 5000c500a5c0e61d SAT3
  sdm
  ✓ 7200 rpm me1111111111111111 zfs_member "chungus" (12.00T) 111111111111119em
  27° no i/o | ST12000NM0007-2A SN02 12.00T 4096B gpt 5000c500b49867e5 SAT3
  [PCI Express 0000:47.00.0 (x16, gen 4.0)]

-----[virtual [0]]-----[-]
  loop2 n/a n/a 0.00 512B none n/a n/a
  press 'v' to dismiss details
  Advanced Micro Devices, Inc. [AMD] FCH SATA Controller [AHCI mode]
  Firmware: F5q BIOS: American Megatrends International, LLC. Load: 12Gbps
  sdm: ST12000NM0007-2ASN03 (10.91TiB) S/N: ZCH0DFDJ WC+ WRV- R0-
  Sectors: 23437770752 (512B logical / 4096B physical) SAT3 (6Gbps)
  Partitioning: gpt I/O scheduler: [mq-deadline] none
  10.91TiB P01 2048-23437752319 sdm1 "zfs-d60b954b8cc5eae0" bf01 1MiB align
  zfs_member "chungus"

  up no i/o | Linux mdadm 1.2 106.23G 512B none root NVMe

```

growlight 1.2.35 (6) Couldn't read link at /sys/class/block/tracefs (No such fil


```

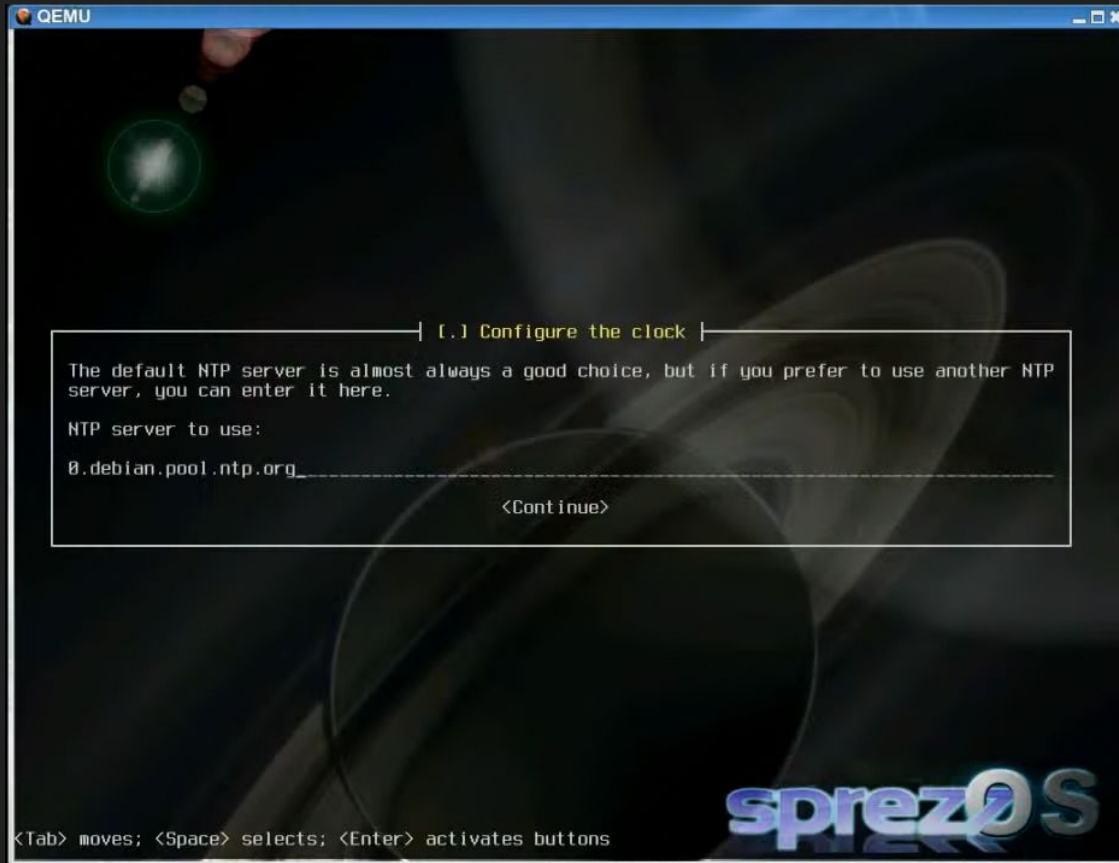
sdd      ST12000NM0007-2A SN02  12.00T 4096B ✓0W△. gpt  5000c500b3f4afb4 SAT3
sr0      iHBS112  2      CL0F   1.07G 512B U0... none n/a          PATA
nvme0n1  WDS100T3X0C-00SJ n/a    1.00T 512B ✓.... gpt  1908E1805012 NVMe
nvme2n1  INTEL MEMPEK1W01 n/a    14.40G 512B ✓.... gpt  PHBT729201SR016D NVMe
nvme1n1  WDS100T3X0C-00SJ n/a    1.00T 512B ✓.... gpt  1908E1801188 NVMe
md127    Linux mdadm  1.2 106.23G 512B VM... none root        NVMe
loop0    n/a        n/a     0.00 512B V0... none n/a          n/a
loop6    n/a        n/a     0.00 512B V0... none n/a          n/a
loop5    n/a        n/a     0.00 512B V0... none n/a          n/a
loop3    n/a        n/a     0.00 512B V0... none n/a          n/a
loop4    n/a        n/a     0.00 512B V0... none n/a          n/a
loop7    n/a        n/a     0.00 512B V0... none n/a          n/a
loop2    n/a        n/a     0.00 512B V0... none n/a          n/a
loop1    n/a        n/a     0.00 512B V0... none n/a          n/a
sdj      STORAGE DEVICE 1203  0.00 512B R0... none n/a          PATA
sdi      STORAGE DEVICE 1203  0.00 512B R0... none n/a          PATA
sdl      STORAGE DEVICE 1203  0.00 512B R0... none n/a          PATA
sdg      STORAGE DEVICE 1203  0.00 512B R0... none n/a          PATA
sdk      STORAGE DEVICE 1203  0.00 512B R0... none n/a          PATA

```

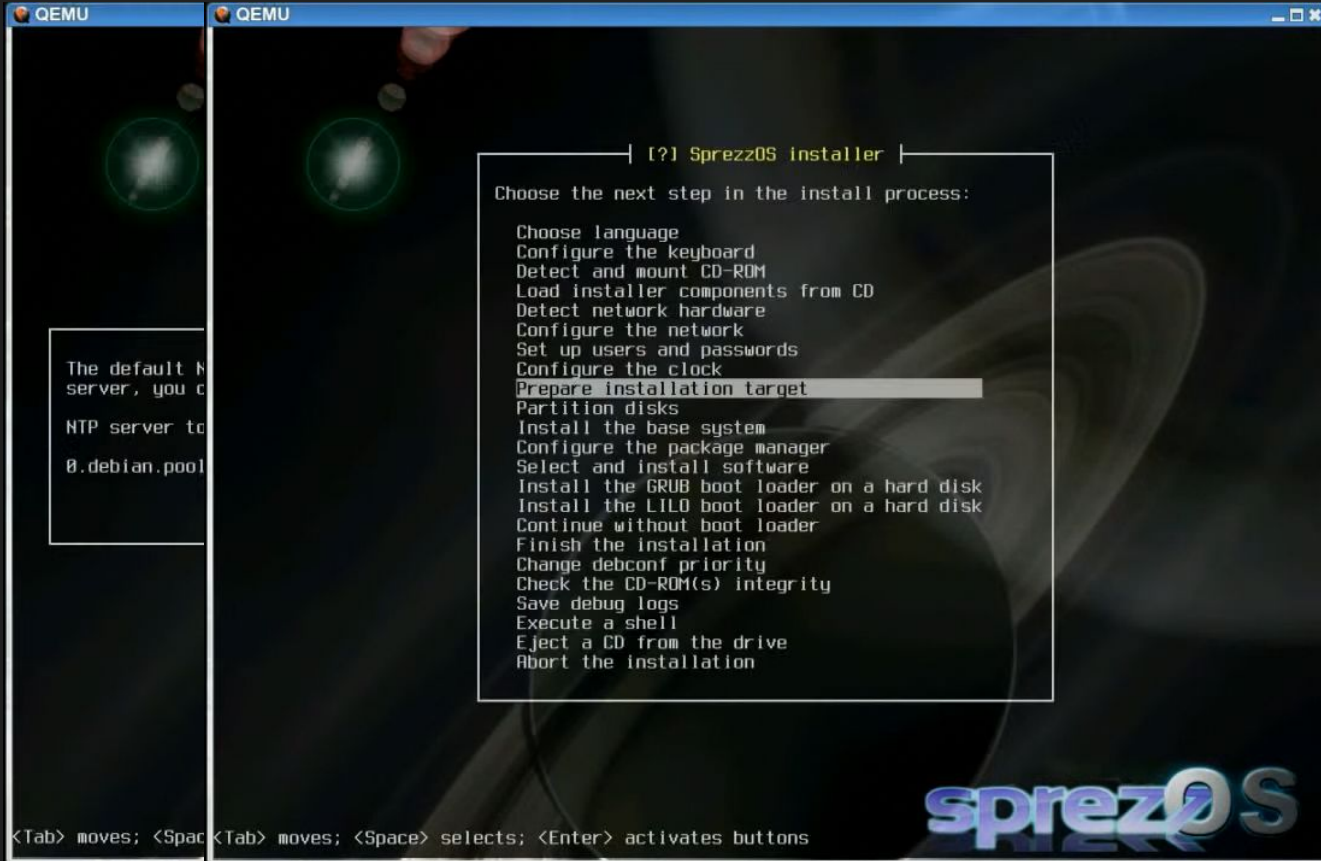
Flags: (R)emovable, (U)nloaded, (V)irtual, (M)dadm, (Z)pool,
(D)M, r(0)tational, (r)ead-only, (W)ritecache enabled,
(B)IOS bootable, v/△: Read-Write-Verify, ✓/x/⊗: SMART status

```
[growlight](0)> █
```

SprezzOS 1.0 installer (d-i fork), 2012



SprezzOS 1.0 installer (d-i fork), 2012



form your parentheses of salt and say the incantation...

in a d-i context:

```
openvt -v -w -s -- \  
fbvfbterm /usr/share/sprezzatech/sprezzos.png \  
growlight-curses -i -t "$TARGET" --disphelp
```

SprezzOS 1.0 installer (d-i fork), 2012

The screenshot shows a QEMU window displaying the SprezzOS 1.0 installer. The main window is titled "QEMU" and shows a disk partitioning screen. The screen is divided into several sections:

- Left Panel:** A vertical menu with options: "Choose the", "Choose l", "Configur", "Detect a", "Load ins", "Detect n", "Configur", "Set up u", "Configur", "Prepare", "Partitio", "Install", "Configur", "Select a", "Install", "Install", "Continue", "Finish t", "Change d", "Check th", "Save deb", "Execute", "Eject a", "Abort th".
- Top Panel:** A list of storage devices with their properties. The "removable" option is highlighted in blue. The list includes:
 - ferromag sdc: 128.84G unpartitioned space, QEMU HARDDISK 1.2, 128.84G 512B none n/a, PATA
 - ferromag sda: 85.89G unpartitioned space, QEMU HARDDISK 1.2, 85.89G 512B none n/a, PATA
 - ferromag sdb: 128.84G unpartitioned space, QEMU HARDDISK 1.2, 128.84G 512B none n/a, PATA
 - removable sr0: Iso9660 "ISOIMAGE" (319.57M) at /cdrom, QEMU DVD-ROM 1.2, 319.57M 2048B dos n/a, PATA
- Bottom Panel:** A list of keyboard shortcuts for navigating and configuring the installer. The list includes:
 - q: quit, ctrl+'L': redraw the screen
 - e: view environment details, 'H': toggle this help display
 - v: view selection details, 'D': view recent diagnostics
 - E: view mounts / targets, 'Z': modify aggregate
 - A: create aggregate, 'Z': destroy aggregate
 - : collapse adapter, '+': expand adapter
 - k/'↑': navigate up, 'j'/'↓': navigate down
 - #PageUp: previous adapter, #PageDown: next adapter
 - /: search, p: configure loop device
 - h/'←': navigate left, l/'→': navigate right
 - m: make partition table, r: remove partition table
 - W: wipe master boot record, 'B': bad blocks check
 - n: new partition, d: delete partition
 - s: set partition attributes, M: make filesystem/swap
 - F: fsck filesystem, w: wipe filesystem
 - U: set filesystem UUID, L: set filesystem label/name
 - o: mount filesystem/swap on, O: unmount filesystem/swap off
 - t: mount target, T: unmount target
 - i: enter target mode, I: leave target mode
 - *: finalize UEFI / '#' finalize BIOS / '@' finalize fstab


At the bottom of the window, there is a footer: "http://wiki.sprezzos.com/dokuwiki/index.php/Download" and "growlflash 1.0.4.1 (0) No ZFS support in this build." The SprezzOS logo is visible in the bottom right corner.

SprezzOS 1.0 installer (d-i fork), 2012

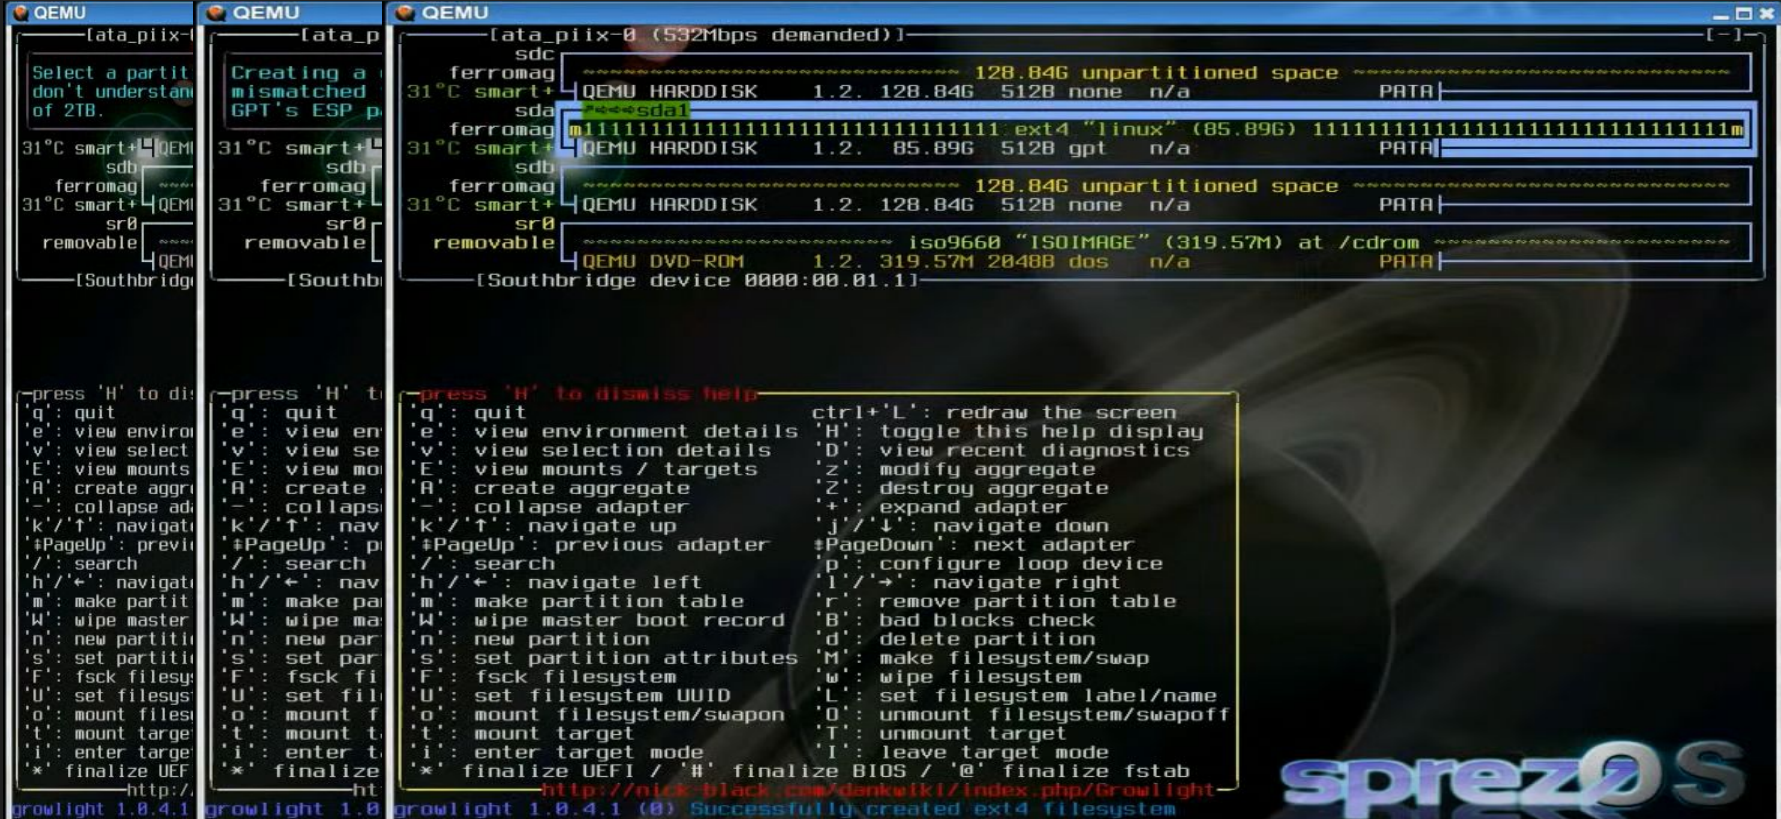
```
QEMU [ata_piix-0 (532Mbps demanded)] [-]-
Select a partition table type. GPT is recommended unless you must use tools and/or hardware which
don't understand it. Please note that MSDOS partition tables do not support partitions in excess
of 2TB.
31°C smart+ QEMU HARDDISK 1.2. 85.89G 512B none n/a
sdb
ferromag 128.84G unpartition
31°C smart+ QEMU HARDDISK 1.2. 128.84G 512B none n/a
sr0
removable iso9660 "ISOIMAGE" (319.
QEMU DVD-ROM 1.2. 319.57M 2048B dos n/a
[Southbridge device 0000:00.01.1]
select a table type
gpt GUID Partition Table
dos IBMPC (DOS) / Master Boot Record
apm Apple Partition Map
Esc returns
press 'H' to dismiss help
q: quit ctrl+'L': redraw the screen
e: view environment details H: toggle this help display
v: view selection details D: view recent diagnostics
E: view mounts / targets Z: modify aggregate
A: create aggregate Z: destroy aggregate
-: collapse adapter +: expand adapter
k / ↑: navigate up j / ↓: navigate down
+PageUp: previous adapter +PageDown: next adapter
/: search p: configure loop device
h / ←: navigate left l / →: navigate right
m: make partition table r: remove partition table
W: wipe master boot record B: bad blocks check
n: new partition d: delete partition
s: set partition attributes M: make filesystem/swap
F: fsck filesystem w: wipe filesystem
U: set filesystem UUID L: set filesystem label/name
o: mount filesystem/swapon O: unmount filesystem/swapoff
t: mount target t: unmount target
i: enter target mode I: leave target mode
* finalize UEFI / # finalize BIOS / @ finalize fstab
http://nick-black.com/dankwiki/index.php/Growlight
growlight 1.8.4.1 (8)
```



SprezzOS 1.0 installer (d-i fork), 2012



SprezzOS 1.0 installer (d-i fork), 2012



Why growlight might be an improvement

- Growlight is an enthusiastically-maintained post-install tool. New hardware and kernel interfaces are quickly supported, and these would become immediately available to the installer. Recent examples:
 - Linux 5.6 "drivetemp"
 - Shingled Magnetic Recording
 - Linux 5.15 `BLKGETDISKSEQ ioctl`
- There's (presumably) less incentive to work on an installer-only component, and a different set of testing opportunities.

Why growlight might be an improvement

- UI naturally supports simultaneous displays of detail for multiple devices
- UI exposes more information (SMART status, firmware versions, etc.)
- Controller-based hierarchy in UI with bandwidth considerations

Partman is arguably the most complex d-i component, *but is restricted to the primitive UI available to the rest of the installer*. This is great for visual continuity, but less great for displaying a lot of information.

Dispelling those concerns which can be dispelled

- UI continuity: change the background color to blue and you're halfway there
- C vs shell: "Partman has a very specific structure and requires a fairly strict conformance to this structure for udebs that extend its functionality." partman is not a trivial system to pick up and extend, even beyond the peculiarities of the d-i environment.
- C vs modern languages. Sorry, this is what I've got--but it's there. The installer would seem to present a rather small attack surface?

Dispelling those concerns which can be dispelled

- Binary size: first-attempt udeb of 1.2.35 runs ~200K (includes both binaries):

```
-rw-r--r-- 1 dank dank 207316 2021-08-16 06:24 growlight_1.2.35-1_amd64.udeb
```

- Also wants a few libraries (libatasmart, libblkid, libpci), among them Notcurses, but only `libnotcurses-core2`, **not** `libnotcurses2`.
- Only the latter has the dependency chains necessary for processing media.
- Probably a wash when `partman-*` and `parted` are removed

- Don't want ZFS? No problem, already built sans-ZFS for the archive.

Dispelling those concerns which can be dispelled

- Can coexist with/fall back to partman, especially at first
- `.isinstallable` udeb control file allows disabling it / removing from menu at runtime for configurations known to be problematic
 - Not recommending this as any long-term solution, of course
- Notcurses ought work fine with serial terminals, over remote console, in pure VGA mode restricted to ASCII, framebuffer, etc.
- "Does it support [insert blockdev thing here]?" Probably. If not, there's almost certainly a bug on it. Might be a flurry of initial development to fill out rough edges.

Concerns remaining concerning

- Derivatives might have their own modifications to partman that would no longer be applicable
- What to do in GUI mode? Could just launch xterm if it's Xorg-based? If Linux framebuffer based, probably sufficient to use `openvt`?
- Existing preseed recipes would need either a bug-compatible reimplementaion of preseeding logic or to keep a copy of partman around

Concerns remaining concerning

- Hasn't seen anywhere near the diversity of testing/validation that partman has...but you've gotta start somewhere.
- Need verify that any replacement works at least as well with screenreaders and braille; I am not qualified to do this
- I "own" growlight; would Debian need its own fork?
 - Could probably make any undesirable aspects config options.
 - I'm actively looking to facilitate its use in the Debian installer; you can rely on a friendly upstream
- No translation work has been done on growlight
 - Could translation investment in partman be easily reappropriated?