

Debian's init system - Past, Present and Future (and some numbers)

Michael Biebl

Linux Plumbers Conference 2010, Boston/MA

4. November 2010

About Debian

- ▶ around 1000 developers
- ▶ around 30000 binary packages
- ▶ base of a lot of derivative distributions
- ▶ flexibility one of its strengths (but also weaknesses)
- ▶ release cycles around 18-24 months
- ▶ strong relationship between package and maintainer (DD/DM), though teams become more and more wide spread

sysvinit / SysV init scripts

- ▶ sysvinit is the default init system in Debian (essential package)
 - 0/6 halt/reboot
- ▶ runlevels: 1 single-user
 - 2-5 multi-user
- ▶ single-user mode semantics are weird (S vs. 1)
- ▶ special runlevel S (/etc/rcS.d/) for system configuration/initialization.
- ▶ /etc/inittab
 - ▶ gettys: 1:2345, 2-6:23
 - ▶ calls /etc/init.d/rc **runlevel**
 - ▶ powerfail and ctrlaltdel
- ▶ interface for package maintainer: invoke-rc.d/update-rc.d
- ▶ start priorities are static, provided by the package maintainer

LSB / insserv

- ▶ LSB header information can be used to calculate the start priorities dynamically, which in turn makes it much easier to fix bugs in the boot sequence
- ▶ dependency information to start services in parallel (startpar)
- ▶ all SysV init scripts were updated to contain an LSB init header
- ▶ effort started with lenny, mostly driven by Petter Reinholdtsen
- ▶ insserv will be enabled by default in squeeze: LSB header information is reasonably well tested
- ▶ LSB logging and helper functions via `/lib/lsb/init-functions` (lsb-base)

LSB header information

Example:

```
### BEGIN INIT INFO
# Provides:          network-manager
# Required-Start:    $remote_fs dbus udev
# Required-Stop:     $remote_fs dbus udev
# Should-Start:      $syslog
# Should-Stop:       $syslog
# Default-Start:     2 3 4 5
# Default-Stop:      0 1 6
# Short-Description: network connection manager
# Description:       Daemon for automatically switching
#                    network connections to the best
#                    available connection.
### END INIT INFO
```

Some statistics...

Analysed LSB header information sorted by popcon¹ data for unstable:

- ▶ 1001 packages shipping 1077 SysV init scripts
- ▶ long tail:
 - ▶ dpkg: 93694 installations
 - ▶ 1. util-linux(93679), 10. acpid (85877), 100. setserial (3348), 500. sleepd (88)
- ▶ only 2 did not have an LSB init header (RC bug)
- ▶ 647 packages sourcing /lib/lsb/init-functions

¹<http://popcon.debian.org>

Some statistics...(contd)

▶ Default-Start:

2 3 4 5	S	"	1 2 3 4 5	S 2 3 4 5	rest
924	119	17	5	3	9

- ▶ 'S': 119, rather long list, needs further inspection
- ▶ " : hm, what? start on shutdown: not supported by systemd
- ▶ 'S 2 3 4 5': portmap/nfs-common: causes problems in systemd

▶ Default-Stop:

0 1 6	"	0 6	1	rest
851	129	52	37	8

- ▶ " : teardown: services which don't need a "clean" shutdown and can simply be killed
- ▶ '0 6' and '1': probably just buggy

current state of systemd in Debian

- ▶ Tollef Fog Heen maintainer of systemd in Debian
- ▶ v11 available from experimental, along with dbus 1.4
- ▶ systemd runs with default kernel from squeeze/sid (cherry-pick for /sys/fs/cgroup patch)
- ▶ uses LSB/sysv init scripts
- ▶ works mostly fine (chokes on nfs-common/portmap though which is part of a default installation unfortunately, dep cycle on shutdown)

integration work / TODO

- ▶ native mount for v12
- ▶ ship more native system unit files in packages (currently on hold due to freeze)
- ▶ hook for lsb-base: redirect `/etc/init.d/foo action` → `systemctl action foo.service`
- ▶ interface for package maintainer scripts, support in debhelper/cdbs
- ▶ clean up `/etc/rcS.d` (tmpfiles as general interface?)
- ▶ kFreeBSD port: figure out, how to support this platform (or not)

looking ahead

- ▶ realistically, we have to plan with at least 3 release cycles (before ripping out any compat code)
- ▶ ship as option in wheezy
- ▶ make it awesome, so users/admins will want to use it
- ▶ providing a fallback plan will make people feel safe to try it
- ▶ allow a gradual transition
- ▶ acknowledge that Debian users are more conservative and expect that upgrades don't break their system
- ▶ there is no mechanism in Debian to force something like systemd into the distribution, but having
 - ▶ a dedicated maintainer/team
 - ▶ a clear view of the impact of such a change
 - ▶ a (rough) consensus among fellow DDs (at least those maintaining relevant packages)
- ▶ makes such a large transition possible