

Beyond Init: Systemd One Year Later

LinuxCon Japan 2011

Lennart Poettering
lennart@poettering.net



June 2011

Who Am I?

Software Engineer at Red Hat, Inc.

Creator of PulseAudio, Avahi, systemd and a few other Free Software projects

<http://0pointer.de/lennart/>

lennart@poettering.net

IRC: mezcadero

One Year Later

First Distribution Based on systemd:

First Distribution Based on systemd:

Fedora 15

2011-05-24

Other distributions following over the next year

We can now boot a system shell-free

We can boot userspace in $<1s$

Goal: Overall boot times of 10s for everybody!

F15: Adoption

F16: Completion

F17: Perfection

Upstream! Upstream! Upstream!

systemd is an Init System

systemd is an Init System
systemd is a Platform

Standardization

De-Balkanization

Integration

Modularization

Bloat?

Bloat? Hell, no!
PID 1 does Unit Control

Bloat? Hell, no!
PID 1 does Unit Control
Basic Set of Auxiliary Services do the rest

Bloat? Hell, no!

PID 1 does Unit Control

Basic Set of Auxiliary Services do the rest

Parallelization, Parallelization, Parallelization

Bloat? Hell, no!

PID 1 does Unit Control

Basic Set of Auxiliary Services do the rest

Parallelization, Parallelization, Parallelization

Embedded Hackers, we love you!

Interfaces

Interfaces

Dracut, udev, D-Bus, Plymouth

Focus: the full Bandwith, like the Linux kernel itself

Focus: the full Bandwith, like the Linux kernel itself
Mobile, Embedded, Desktop, Server

Use in Embedded: Yocto, Angstrom

Koen's Blog

F16 Plans: Focus on user sessions and multi-seat

<http://lwn.net/Articles/441328/>

“systemd is a system and session manager for Linux,

“systemd is a system and session manager for Linux, compatible with SysV and LSB init scripts.

“systemd is a system and session manager for Linux, compatible with SysV and LSB init scripts. systemd provides aggressive parallelization capabilities,

“systemd is a system and session manager for Linux, compatible with SysV and LSB init scripts. systemd provides aggressive parallelization capabilities, uses socket and D-Bus activation for starting services,

“systemd is a system and session manager for Linux, compatible with SysV and LSB init scripts. systemd provides aggressive parallelization capabilities, uses socket and D-Bus activation for starting services, offers on-demand starting of daemons,

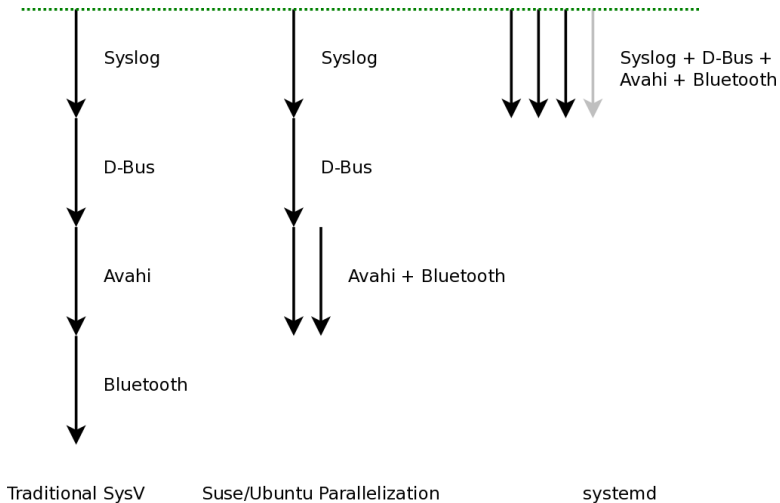
“systemd is a system and session manager for Linux, compatible with SysV and LSB init scripts. systemd provides aggressive parallelization capabilities, uses socket and D-Bus activation for starting services, offers on-demand starting of daemons, keeps track of processes using Linux cgroups,

“systemd is a system and session manager for Linux, compatible with SysV and LSB init scripts. systemd provides aggressive parallelization capabilities, uses socket and D-Bus activation for starting services, offers on-demand starting of daemons, keeps track of processes using Linux cgroups, supports snapshotting and restoring of the system state,

“systemd is a system and session manager for Linux, compatible with SysV and LSB init scripts. systemd provides aggressive parallelization capabilities, uses socket and D-Bus activation for starting services, offers on-demand starting of daemons, keeps track of processes using Linux cgroups, supports snapshotting and restoring of the system state, maintains mount and automount points

“systemd is a system and session manager for Linux, compatible with SysV and LSB init scripts. systemd provides aggressive parallelization capabilities, uses socket and D-Bus activation for starting services, offers on-demand starting of daemons, keeps track of processes using Linux cgroups, supports snapshotting and restoring of the system state, maintains mount and automount points and implements an elaborate transactional dependency-based service control logic.

“systemd is a system and session manager for Linux, compatible with SysV and LSB init scripts. systemd provides aggressive parallelization capabilities, uses socket and D-Bus activation for starting services, offers on-demand starting of daemons, keeps track of processes using Linux cgroups, supports snapshotting and restoring of the system state, maintains mount and automount points and implements an elaborate transactional dependency-based service control logic. It can work as a drop-in replacement for sysvinit.”



Socket-Based Activation

Socket-Based Activation

The kernel orders and buffers requests for us!

Socket-Based Activation

The kernel orders and buffers requests for us!

Implicit dependencies!

Socket-Based Activation

The kernel orders and buffers requests for us!

Implicit dependencies!

Patching daemons

Bus-Based Activation

Starting Less: On-Demand Loading

Parallelizing File System Jobs

Parallelizing File System Jobs autofs!

Shell is evil

Shell is evil

Move to systemd, daemons, kernel, udev, ...

Shell is evil

Move to systemd, daemons, kernel, udev, . . .

Provide proper debugging facilities

The best babysitter.

The best babysitter.
Control Groups!

The best babysitter II

The best babysitter II

Environment, resource limits, working directory, `chroot()`, `umask`, OOM adjustment, nice level, IO priority and class, CPU scheduler priority and policy/reset-on-fork, CPU affinity, timer slack, `stdio` to `syslog/tty/null/kmsg`, `uid`, `gid`, supplementary groups, file system namespacing (r/o file systems, inaccessible systems, mount propagation, private `/tmp`), capabilities (inherited set, bounding set, secure bits), ...

Unit types: service, socket, device, mount, automount, target,
snapshot, timer, swap, path

Don't reinvent the wheel:

Don't reinvent the wheel:

Read SysV/LSB init script headers, read `/etc/fstab`, support traditional `inetd` modes, support `/dev/initctl`, `utmp`, `wtmp`, support `double-fork()`ing daemons.

Don't reinvent the wheel:

Read SysV/LSB init script headers, read `/etc/fstab`, support traditional `inetd` modes, support `/dev/initctl`, `utmp`, `wtmp`, support `double-fork()`ing daemons.

`.desktop` files.

Snapshots

Transaction System

D-Bus!

Substantial coverage of basic OS boot-up tasks,

Substantial coverage of basic OS boot-up tasks, including fsck,

Substantial coverage of basic OS boot-up tasks, including fsck, mount,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock, readahead,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock, readahead, tmpfiles,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock, readahead, tmpfiles, random-seed,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock, readahead, tmpfiles, random-seed, console,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock, readahead, tmpfiles, random-seed, console, static module loading,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock, readahead, tmpfiles, random-seed, console, static module loading, early syslog,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock, readahead, tmpfiles, random-seed, console, static module loading, early syslog, plymouth,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock, readahead, tmpfiles, random-seed, console, static module loading, early syslog, plymouth, shutdown,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock, readahead, tmpfiles, random-seed, console, static module loading, early syslog, plymouth, shutdown, kexec,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock, readahead, tmpfiles, random-seed, console, static module loading, early syslog, plymouth, shutdown, kexec, SELinux,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock, readahead, tmpfiles, random-seed, console, static module loading, early syslog, plymouth, shutdown, kexec, SELinux, initrd+initrd-less boots,

Substantial coverage of basic OS boot-up tasks, including fsck, mount, quota, hwclock, readahead, tmpfiles, random-seed, console, static module loading, early syslog, plymouth, shutdown, kexec, SELinux, initrd+initrd-less boots, cryptsetup, . . .

systemadm

Say No! to Copyright Assignment.

That's all, folks.

That's all, folks.
Any questions?

systemd

<http://www.freedesktop.org/wiki/Software/systemd>

<http://0pointer.de/blog/projects/systemd>

<git://anongit.freedesktop.org/systemd>

#systemd on irc.freenode.org