Comparing systemd and other init processes for embedded systems

Overview about historic and present concepts of UNIX init processes

Jan Altenberg
Open Source Automation Development Lab (OSADL) eG

COMPACT
OSADL
ONLINE
LECTURES

The boot process

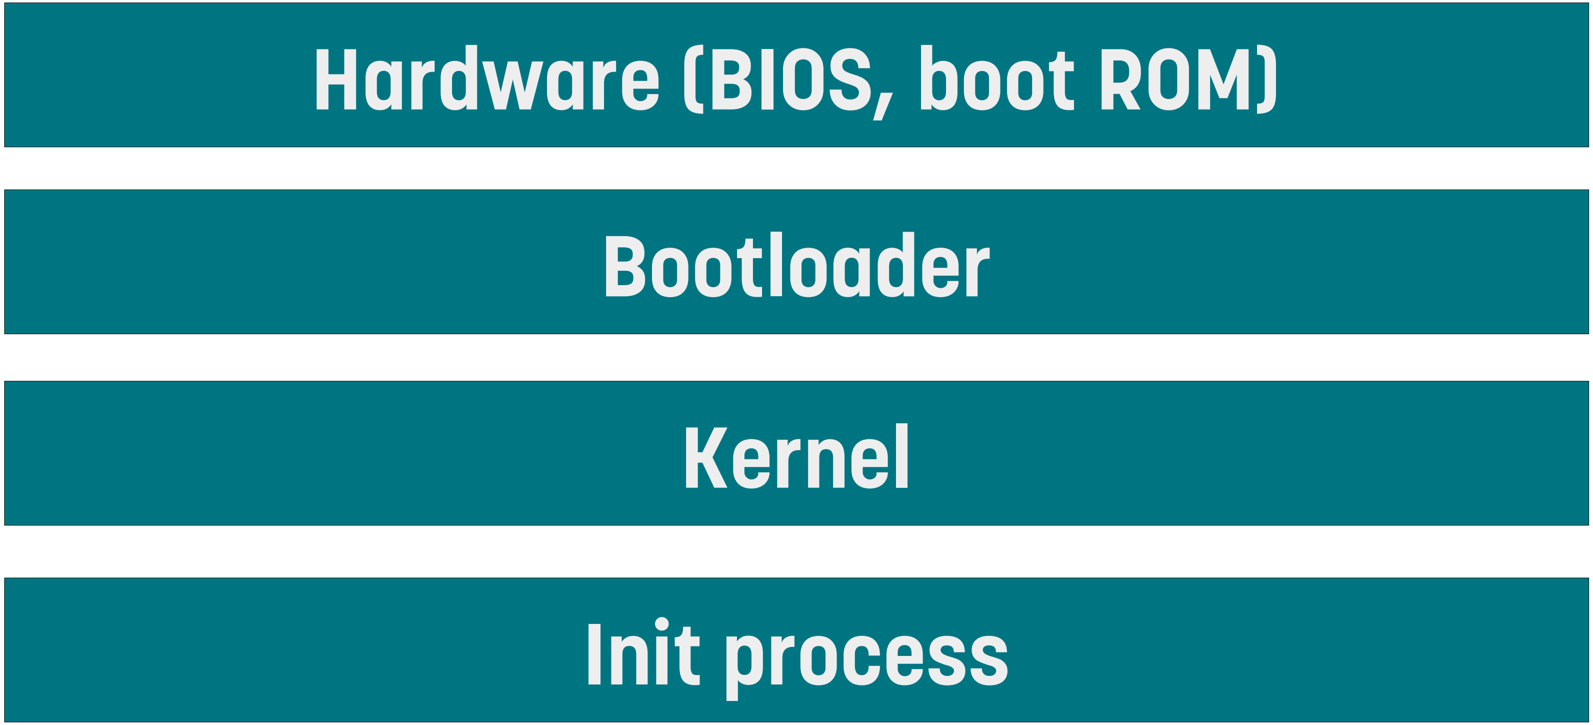

Hardware (BIOS, boot ROM)

Bootloader

Kernel

Init process

PID 1

linux/init/main.c:

```
if (!try_to_run_init_process("/sbin/init") ||
    !try_to_run_init_process("/etc/init") ||
    !try_to_run_init_process("/bin/init") ||
    !try_to_run_init_process("/bin/sh"))
        return 0;
```

PID 1

```
# dpkg -S /sbin/init
systemd-sysv: /sbin/init
```

PID 1

```
# dpkg -S /sbin/init
systemd-sysv: /sbin/init
```

```
# ls -l /sbin/init
/sbin/init ->
/lib/systemd/systemd
```

PID 1

```
# dpkg -S /sbin/init
systemd-sysv: /sbin/init
```

```
# dpkg -S /sbin/init
upstart: /sbin/init
```

PID 1

```
# dpkg -S /sbin/init
systemd-sysv: /sbin/init

# dpkg -S /sbin/init
upstart: /sbin/init

# dpkg -S /sbin/init
sysvinit: /sbin/init
```

History

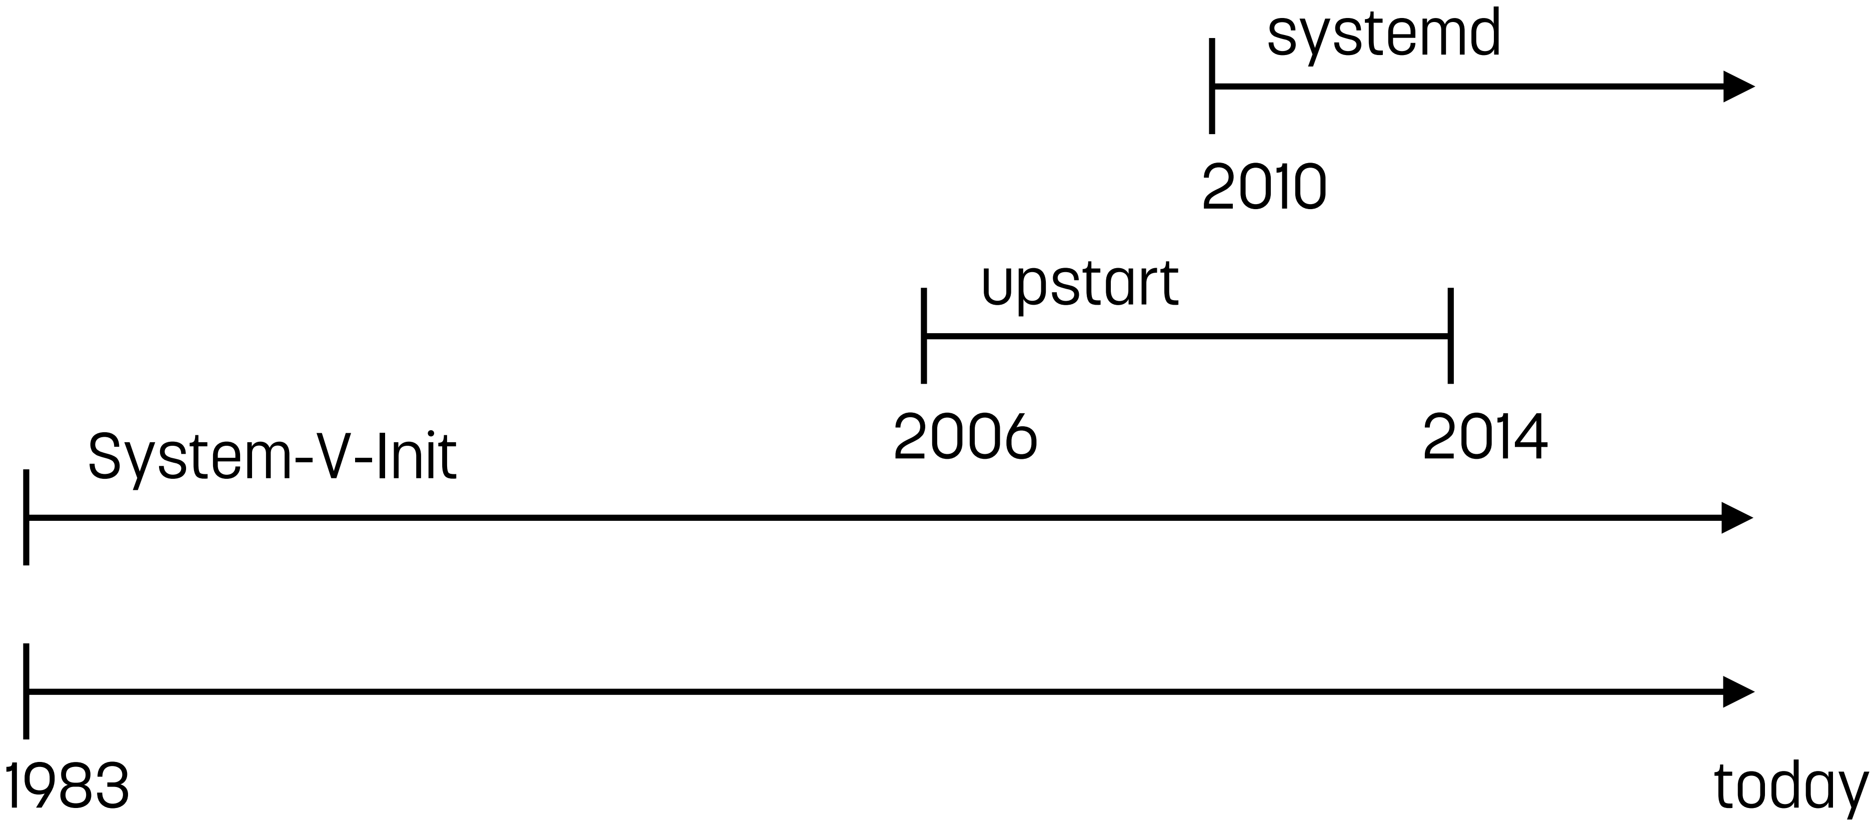

systemd

2010

upstart

2006

2014

System-V-Init

1983

today

Relevant kernel parameters

- init=
  Run a specified binary as init process

- rdinit=
  Run a specified binary as init process for a ramdisk

Initramfs and switch_root

```
# This script has to be run as PID 1!
# 1) Do some work (e.g. drawing a splashscreen)
# 2) Mount the “production system”
mount /dev/my_root_device /production_system

# 3) Unmount filesystems and move the virtual filesystems over
# to the production system
mount --move /dev  /production_system/dev
mount --move /proc /production_system/proc
mount --move /sys  /production_system/sys

# 4) Hand over to the production system using switch_root
# NOTE: Always run switch_root with exec, because it has to run as PID 1!
exec switch_root /production_system /sbin/init
```

System-V-Init

- Originally developed for UNIX System V in 1983
- Available on different UNIX flavors and clones
- Running scripts in sequential order
- System states are organized in runlevels
- Very low number of runtime dependencies

System-V-Init: Runlevels

Start / stop scripts are located in /etc/init.d:

```
/etc/init.d/atd
/etc/init.d/bootclean
/etc/init.d/bootlogd
/etc/init.d/cron
/etc/init.d/exim4
/etc/init.d/glibc.sh
/etc/init.d/gpm
[...]
```

System-V-Init: Runlevels

Start / stop scripts are located in /etc/init.d:

```
/etc/init.d/atd
/etc/init.d/bootclean
/etc/init.d/bootlogd
/etc/init.d/cron
/etc/init.d/exim4
/etc/init.d/glibc.sh
/etc/init.d/gpm
[...]
```

start	Start the daemon
stop	Stop the daemon
status	Check if the daemon is running
reload	Reload configuration
force-reload	Reload configuration (possibly with reboot)

System-V-Init: Runlevels

Runlevels are defined in /etc/rcX.d/

```
K20exim4 -> ../init.d/exim4
S10sysklogd -> ../init.d/sysklogd
S11klogd -> ../init.d/klogd
S18portmap -> ../init.d/portmap
S20gpm -> ../init.d/gpm
S20lldpd -> ../init.d/lldpd
S20lprng -> ../init.d/lprng
[...]
```

System-V-Init: Runlevels

Runlevels are defined in /etc/rcX.d/

```
K20exim4 -> ../init.d/exim4
S10sysklogd -> ../init.d/sysklogd
S11klogd -> ../init.d/klogd
S18portmap -> ../init.d/portmap
S20gpm -> ../init.d/gpm
S20lldpd -> ../init.d/lldpd
S20lprng -> ../init.d/lprng
[...]
```

/etc/init.d/exim4 stop

System-V-Init: Runlevels

Runlevels are defined in /etc/rcX.d/

```
K20exim4 -> ../init.d/exim4
S10sysklogd -> ../init.d/sysklogd
S11klogd -> ../init.d/klogd
S18portmap -> ../init.d/portmap
S20gpm -> ../init.d/gpm
S20lldpd -> ../init.d/lldpd
S20lprng -> ../init.d/lprng
[...]
```

/etc/init.d/klogd start

System-V-Init: Runlevels

0	Shutdown
1	Single user without network
2	Multi user without network
3	Multi user with network
5	Multi user with network and graphical interface
6	Reboot

System-V-Init: /etc/inittab

```
2:2345:respawn:/sbin/mingetty tty2
3:2345:respawn:/sbin/mingetty tty3
```

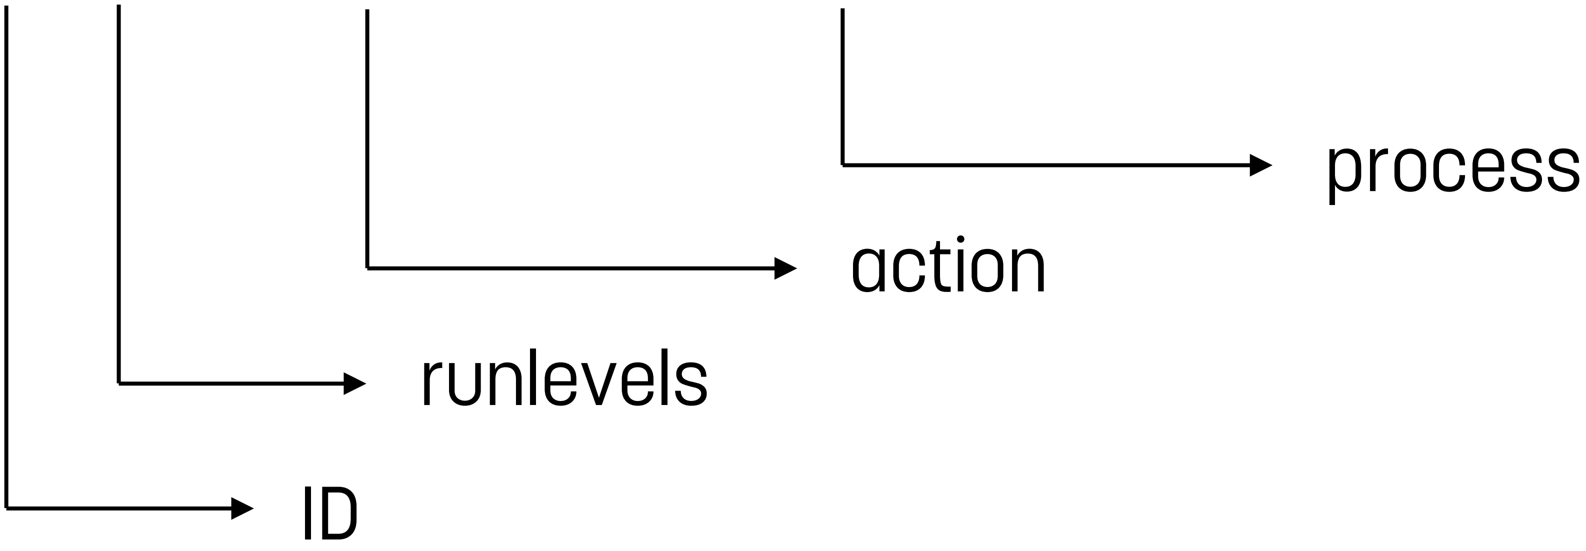

System-V-Init: /etc/inittab

boot	Start the process independently to the runlevels
bootwait	Equal to “boot”. Waits until the process terminates
once	Start once when reaching a relevant runlevel
respawn	Restart the process if terminated
wait	Equal to “once”. Waits until the process terminates

System-V-Init: /etc/inittab

```
l0:0:wait:/etc/init.d/rc 0
l1:1:wait:/etc/init.d/rc 1
l2:2:wait:/etc/init.d/rc 2
l3:3:wait:/etc/init.d/rc 3
l5:5:wait:/etc/init.d/rc 5
l6:6:wait:/etc/init.d/rc 6
```

System-V-Init: LSB headers

```
#! /bin/sh
### BEGIN INIT INFO
# Provides:          lldpd
# Required-Start:    $remote_fs $network $syslog
# Required-Stop:     $network $remote_fs $syslog
# Default-Start:     2 3 4 5
# Default-Stop:      0 1 6
# Short-Description: LLDP daemon
# Description:       lldpd is a 802.1AB implementation, a L2 network
#                    discovery protocol. It also supports CDP, EDP and
#                    various other protocols.
### END INIT INFO
```

System-V-Init: LSB headers

```
#! /bin/sh
### BEGIN INIT INFO
# Provides:          lldp
# Required-Start:    $remote_fs $ne
# Required-Stop:     $network $r        _fs $syslog
# Default-Start:     2 3 4 5
# Default-Stop:      0 1 6
# Short-Description: LLDP daemon
# Description:       lldpd is a 802.1AB implementation, a L2 network
#                    discovery protocol. It also supports CDP, EDP and
#                    various other protocols.
### END INIT INFO
```

Symbolic links for runlevels can be automatically updated, e.g. with `update-rc.d`

System-V-Init: bootchartd

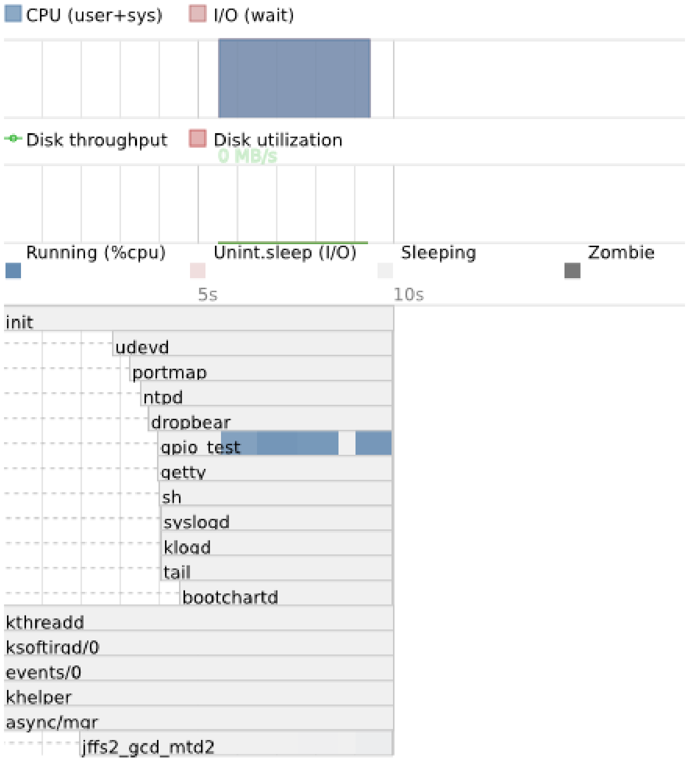

System-V-Init: Is it still relevant?

System-V-Init: Is it still relevant?

YES!

upstart

- First release in 2006
- Mostly used on Ubuntu Linux
- Development is discontinued
  - Final release in 2014

upstart

- Backward compatibility with System-V-Init
- Event based, operates asynchronously
- Supervision of processes while they are running
- Limited number of runtime dependencies

upstart: /etc/init/*.conf

```
# networking.conf
description     "configure virtual network devices"

start on (local-filesystems
          and stopped udevtrigger)

task

pre-start exec mkdir -p /var/run/network

exec ifup -a
```

upstart: /etc/init/*.conf

```
# networking.conf
description     "configure virtual network devices"

start on (local-filesystems
          and stopped udevtrigger)

task

pre-start exec mkdir -p /var/run/network

exec ifup -a
```

dependencies

upstart: /etc/init/*.conf

```
# networking.conf
description     "configure virtual network devices"

start on (local-filesystems
          and stopped udevtrigger)

task

pre-start exec mkdir -p /var/run/network

exec ifup -a
```

Type

service: Daemon

task: Short running process

upstart: /etc/init/*.conf

```
# networking.conf
description     "configure virtual network devices"

start on (local-filesystems
          and stopped udevtrigger)

task

pre-start exec mkdir -p /var/run/network

exec ifup -a
```

Execute BEFORE the process will be run

upstart: /etc/init/*.conf

```
# networking.conf
description     "configure virtual network devices"

start on (local-filesystems
          and stopped udevtrigger)

task

pre-start exec mkdir -p /var/run/network

exec ifup -a
```

Process to be executed

systemd

- First release in 2010
- Announced in the blog post "Rethinking PID 1"
  - http://0pointer.de/blog/projects/systemd.html
- Inspired by launchd
- Not just a single daemon: It provides a whole <u>software collection for system and service management</u>
- High number of runtime dependencies

systemd

- Backward compatibility with System-V-Init
  - Supports System-V and LSB init scripts
- Dependency handling / operates asynchronously
- Supervision of processes while they are running
- Additional services:
  - Login management
  - Network management
  - Watchdog
  - …

systemd: Overview

kernel

CGROUPS

timerfd

...

systemd: Overview

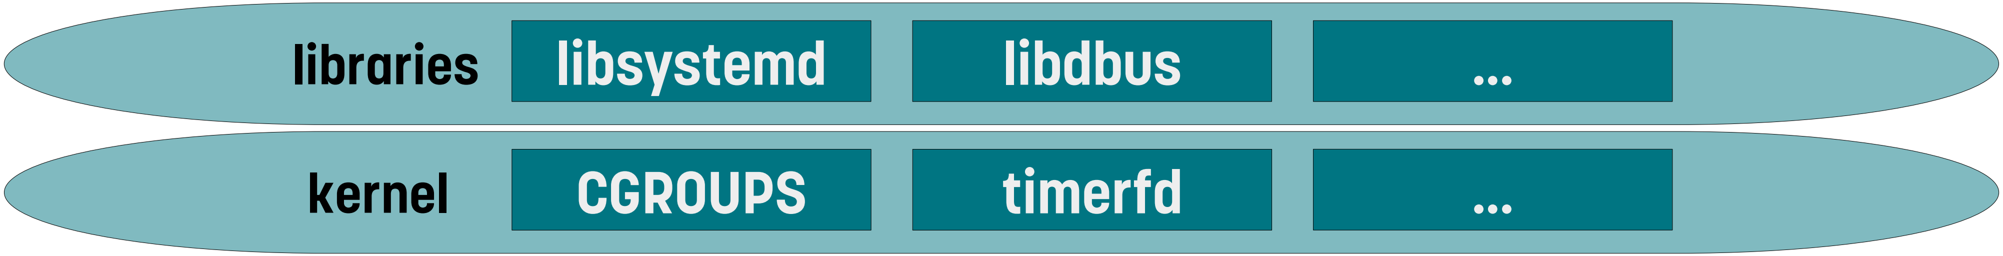

libraries

libsystemd

libdbus

...

kernel

CGROUPS

timerfd

...

systemd: Dependencies

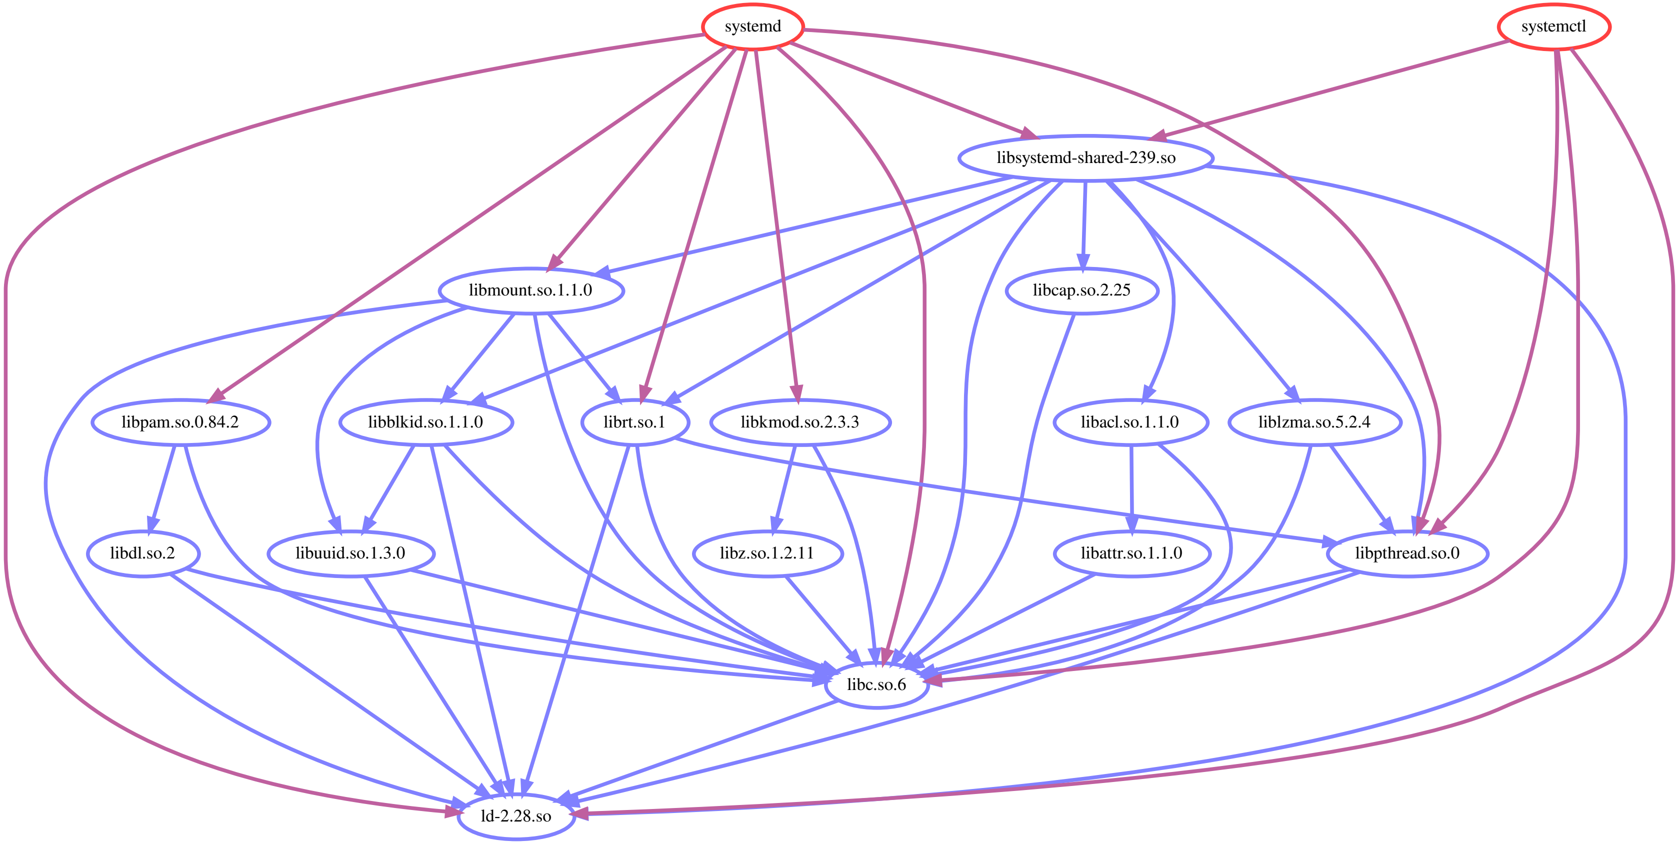

systemd: Overview

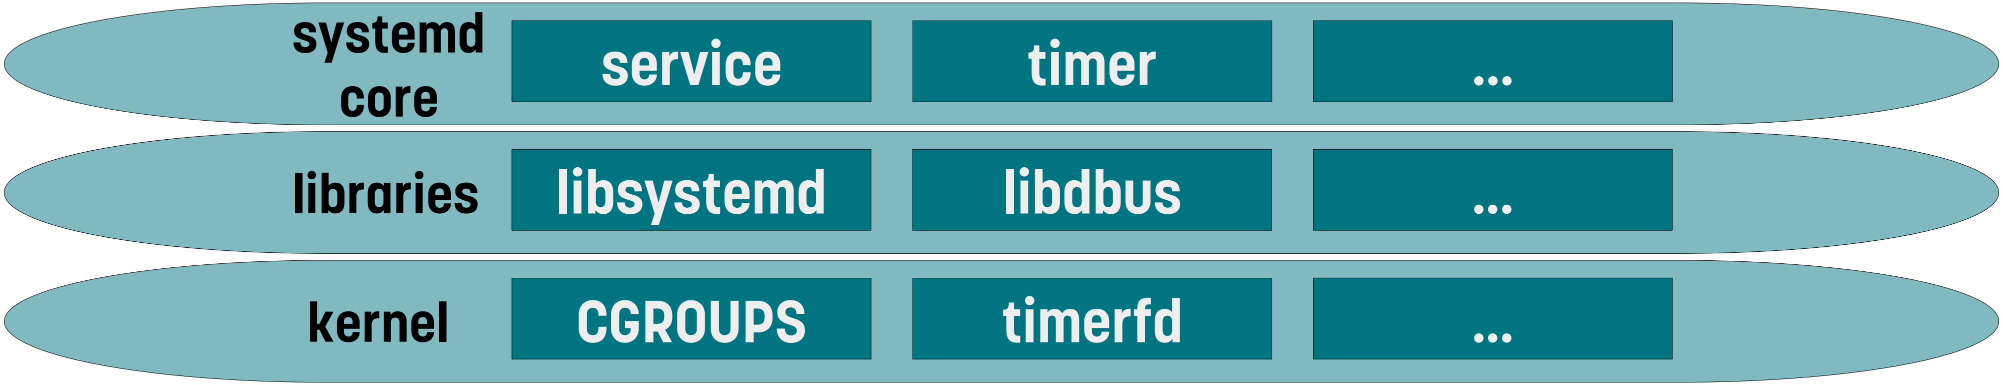

systemd core	service	timer	...
libraries	libsystemd	libdbus	...
kernel	CGROUPS	timerfd	...

systemd: Overview

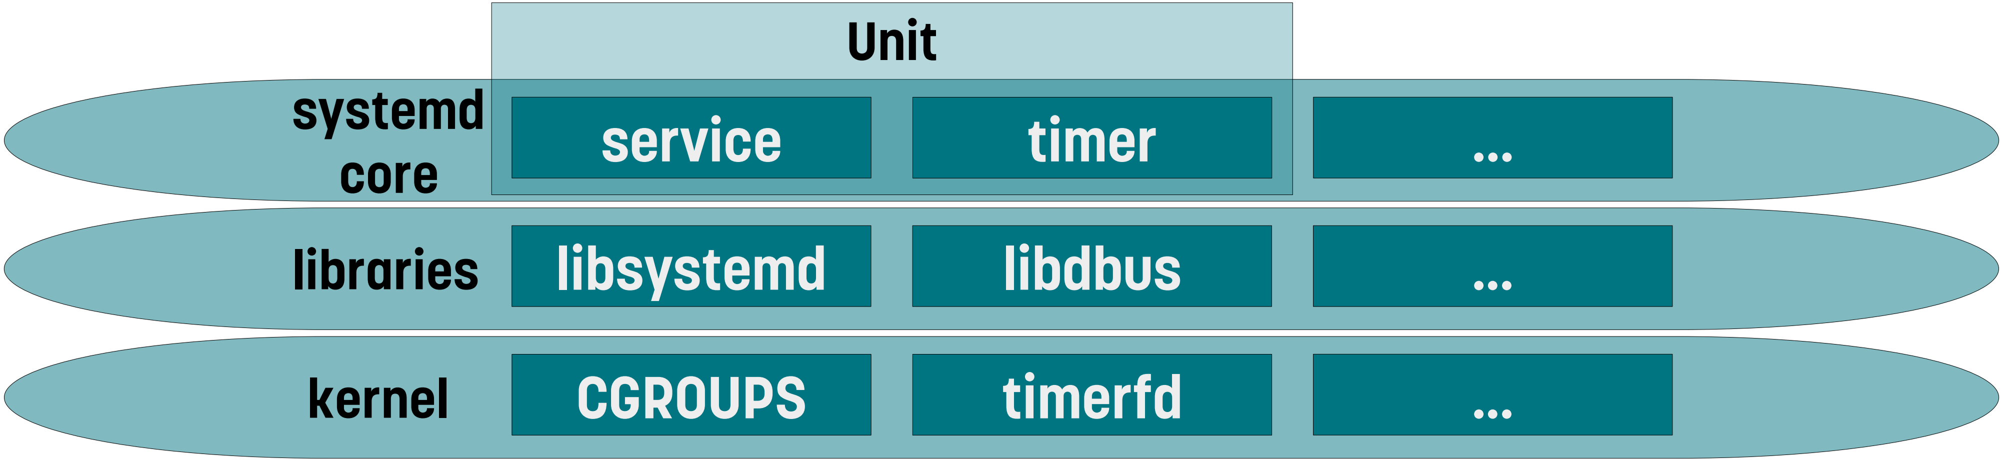

systemd: Overview

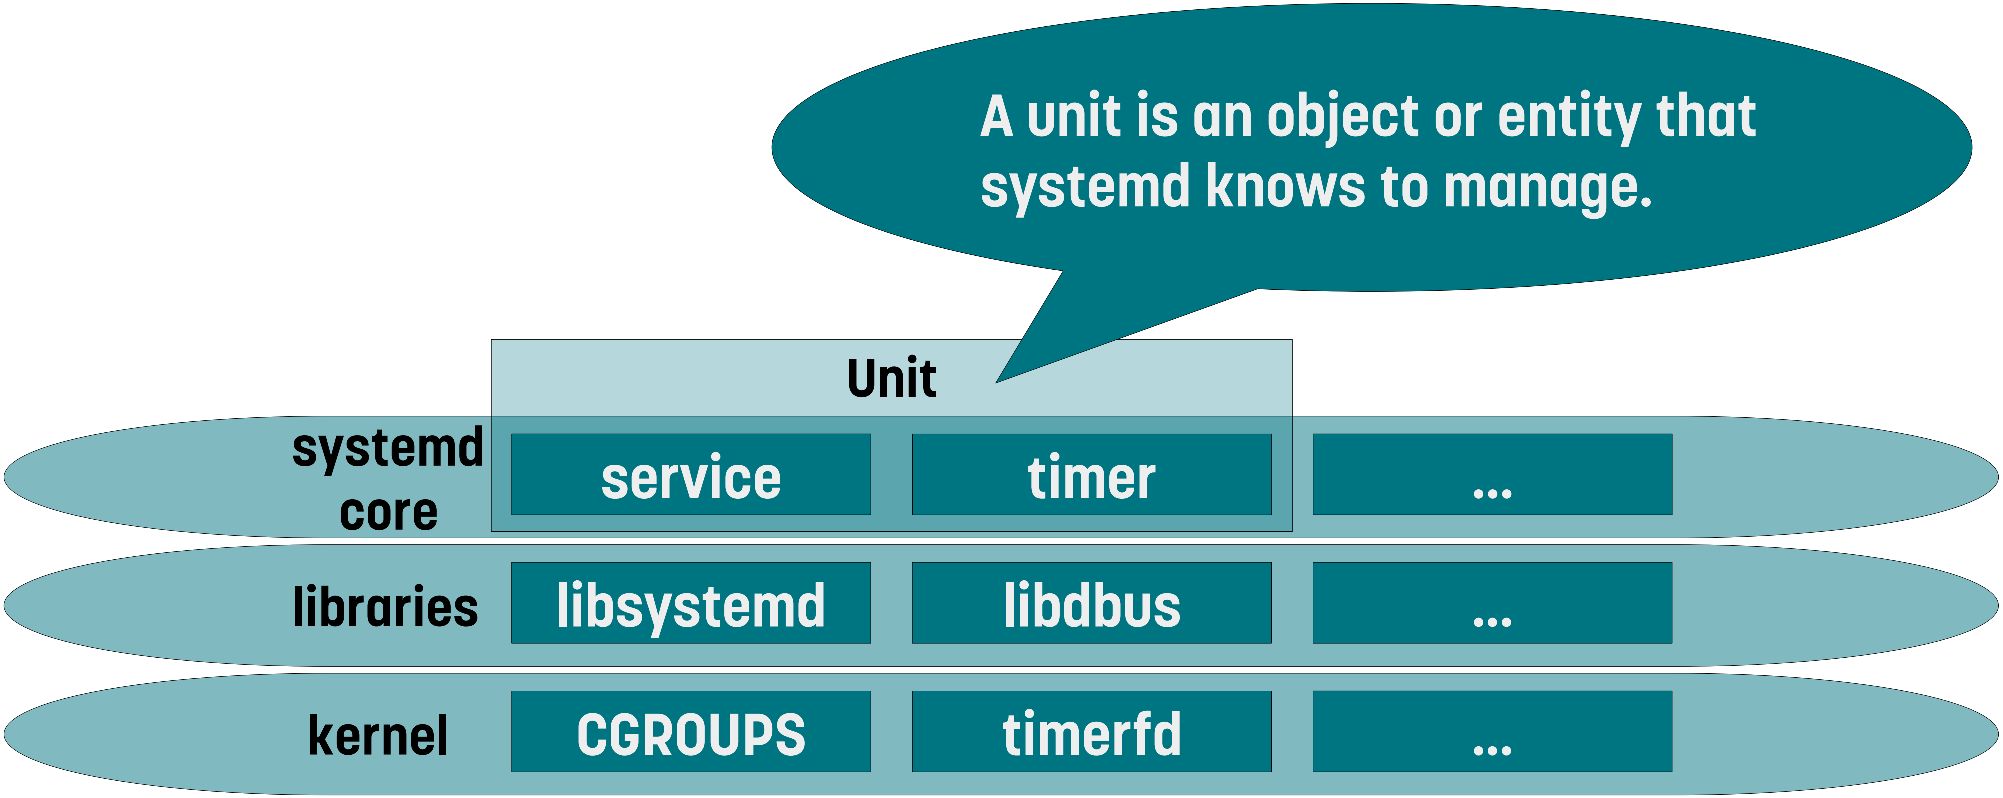

systemd: Overview

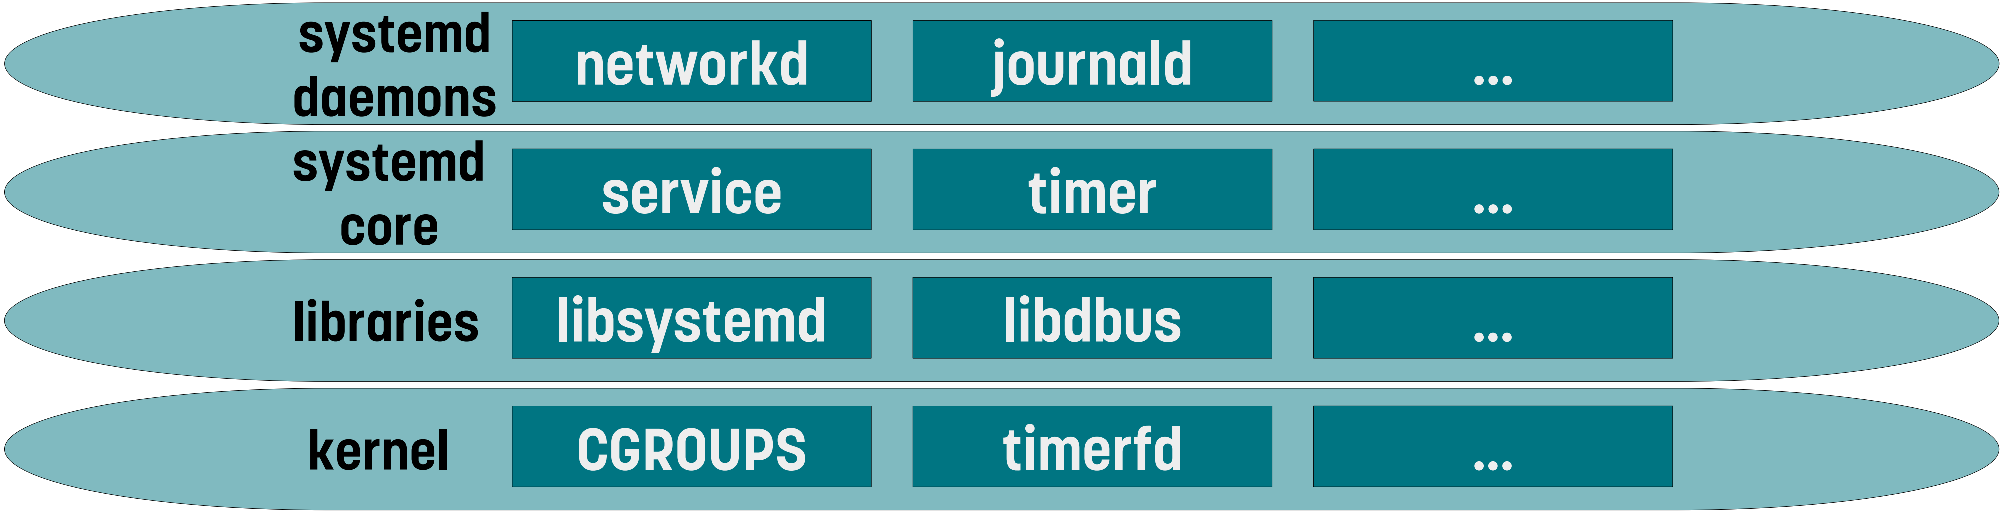

systemd daemons	networkd	journald	...
systemd core	service	timer	...
libraries	libsystemd	libdbus	...
kernel	CGROUPS	timerfd	...

systemd: Overview

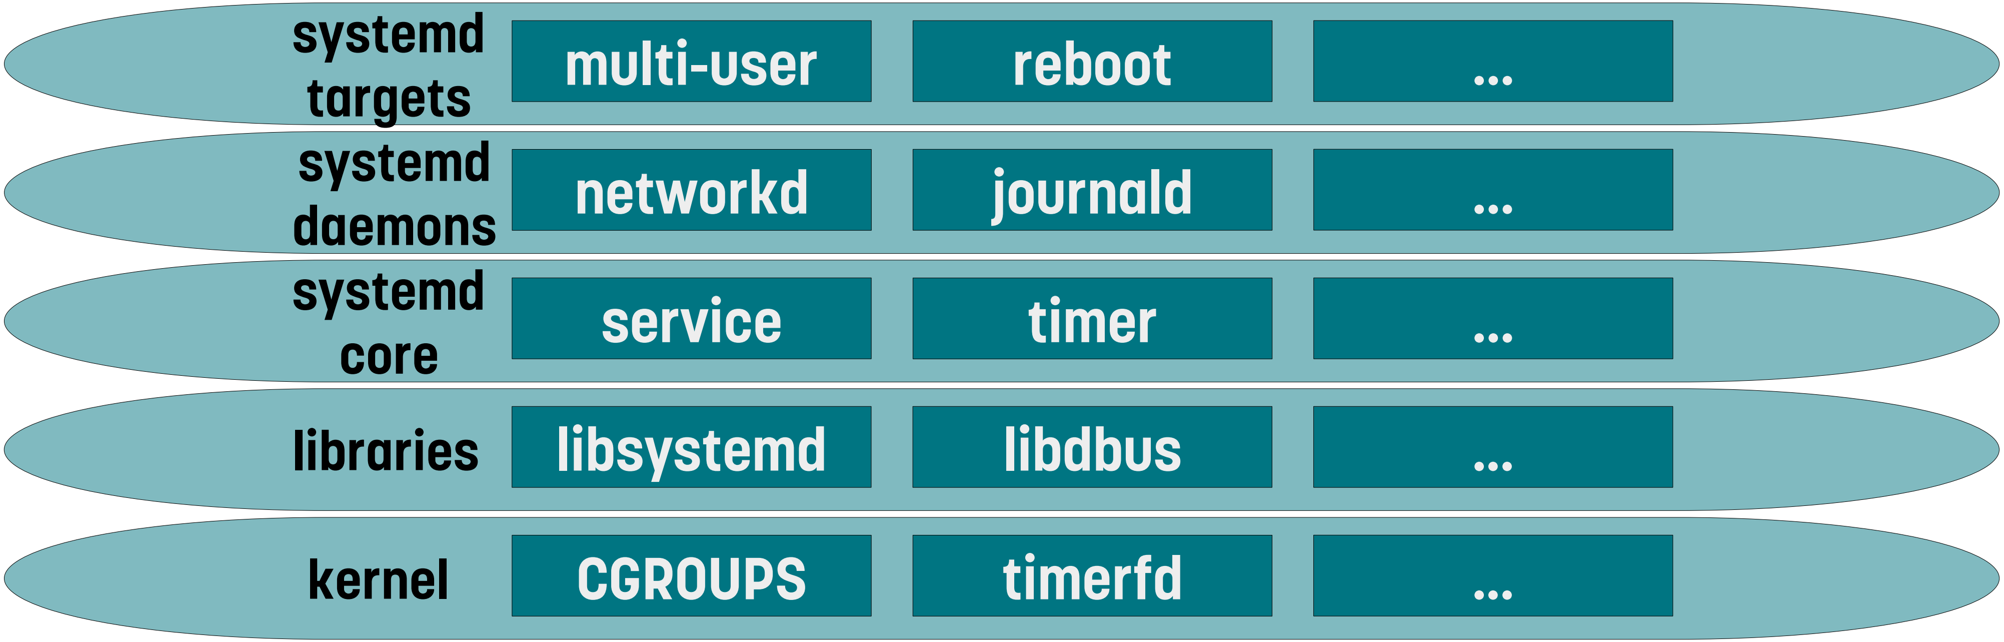

systemd targets	multi-user	reboot	...
systemd daemons	networkd	journald	...
systemd core	service	timer	...
libraries	libsystemd	libdbus	...
kernel	CGROUPS	timerfd	...

systemd: Overview

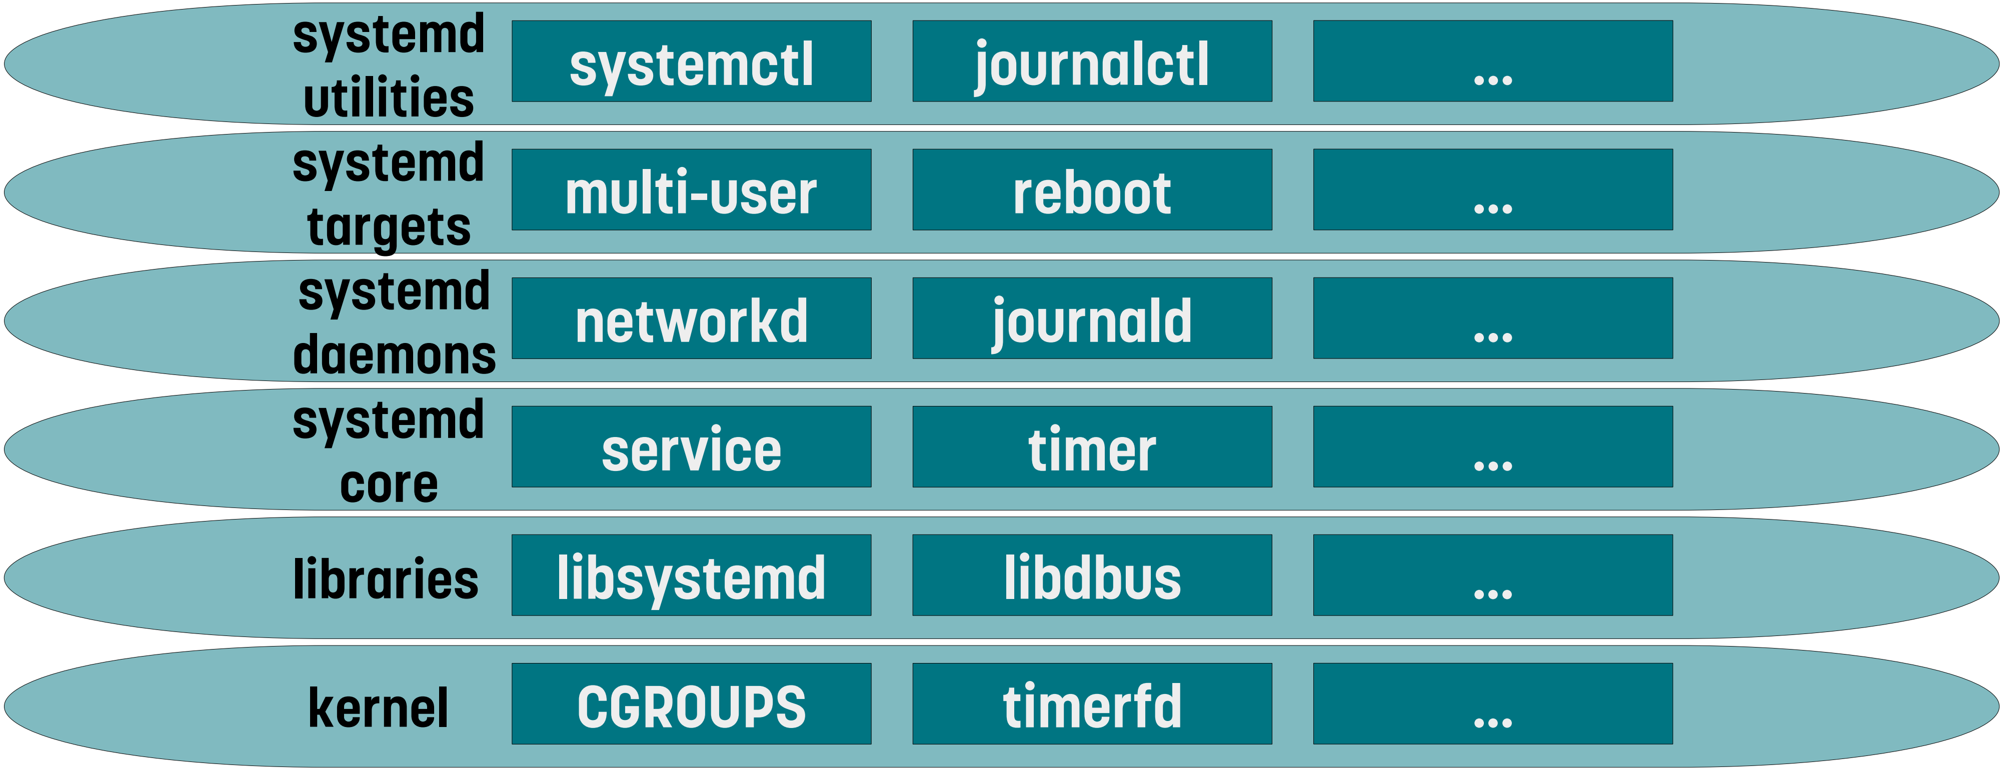

systemd utilities	systemctl	journalctl	...
systemd targets	multi-user	reboot	...
systemd daemons	networkd	journald	...
systemd core	service	timer	...
libraries	libsystemd	libdbus	...
kernel	CGROUPS	timerfd	...

systemd: Targets

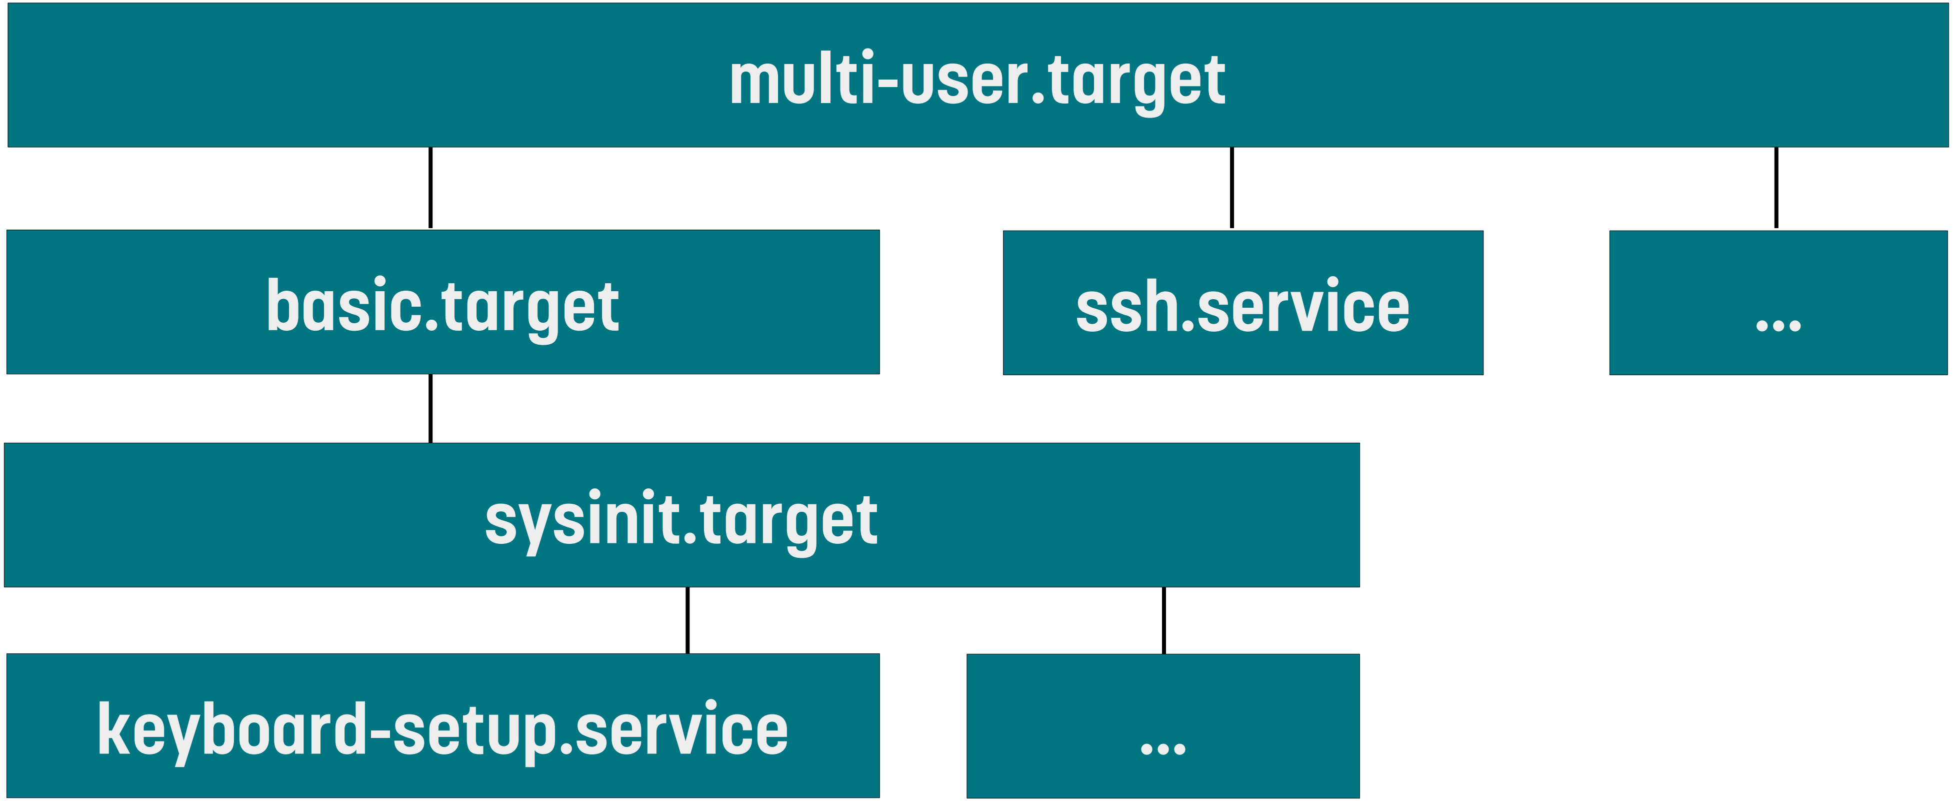

multi-user.target

basic.target

ssh.service

...

sysinit.target

keyboard-setup.service

...

systemd: Targets

Target	System-V runlevel	Description
rescue.target	1	Basic system with rescue shell
multi-user.target	3	All services are started, but no graphical interface is started
graphical.target	5	multi-user.target + graphical interface
reboot.target	6	System reboot

systemd: Controversies

- Build your own opinion!
- Some historical resources:
  - https://lwn.net/Articles/572805/
  - https://www.phoronix.com/news/MTU1NjA
  - https://lists.debian.org/debian-ctte/2014/11/msg00071.html
  - https://people.debian.org/~stapelberg/docs/systemd-dependencies.html

Summary

	Complexity	Number of runtime dependencies	Maintained	Runnable on non-Linux systems
System-V-Init	low	very low	yes	yes
upstart	medium	low – medium	no	probably yes
systemd	medium-high	high	yes	no

Summary

	Dependency management	Parallel execution	Supervision	Built-in additional services
System-V-Init	limited	no	no	no
upstart	yes	yes	yes	no
systemd	yes	yes	yes	yes