

NAME

bootup - System bootup process

DESCRIPTION

A number of different components are involved in the system boot. Immediately after power-up, the system BIOS will do minimal hardware initialization, and hand control over to a boot loader stored on a persistent storage device. This boot loader will then invoke an OS kernel from disk (or the network). In the Linux case, this kernel (optionally) extracts and executes an initial RAM disk image (initrd), such as generated by **dracut(8)**, which looks for the root file system (possibly using **systemd(1)** for this). After the root file system is found and mounted, the initrd hands over control to the host's system manager (such as **systemd(1)**) stored on the OS image, which is then responsible for probing all remaining hardware, mounting all necessary file systems and spawning all configured services.

On shutdown, the system manager stops all services, unmounts all file systems (detaching the storage technologies backing them), and then (optionally) jumps back into the initrd code which unmounts/detaches the root file system and the storage it resides on. As a last step, the system is powered down.

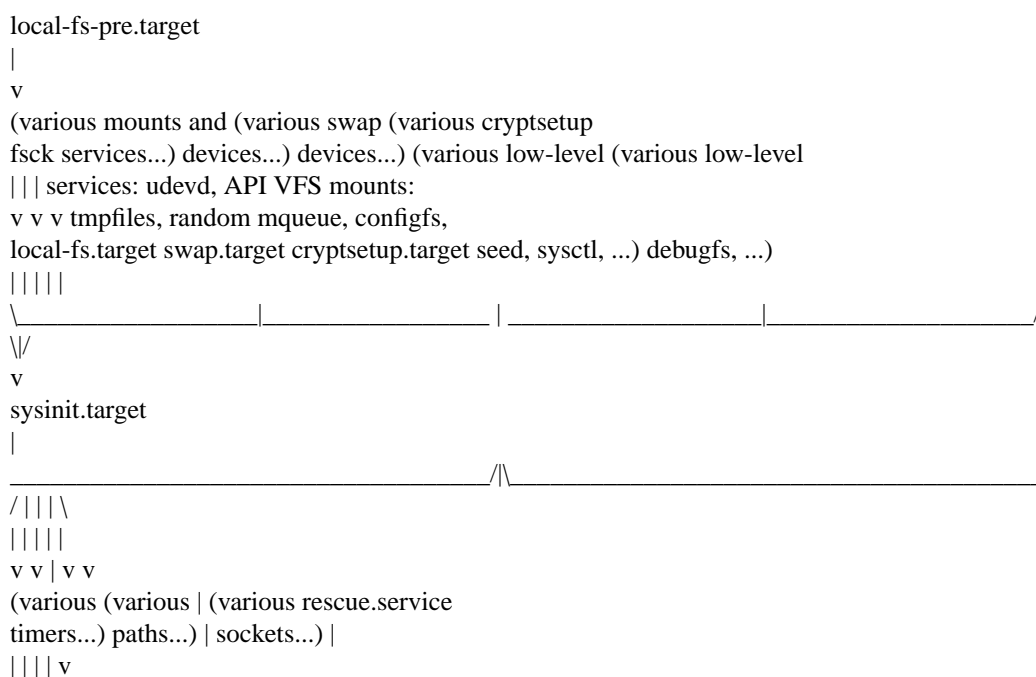
Additional information about the system boot process may be found in **boot(7)**.

SYSTEM MANAGER BOOTUP

At boot, the system manager on the OS image is responsible for initializing the required file systems, services and drivers that are necessary for operation of the system. On **systemd(1)** systems, this process is split up in various discrete steps which are exposed as target units. (See **systemd.target(5)** for detailed information about target units.) The boot-up process is highly parallelized so that the order in which specific target units are reached is not deterministic, but still adheres to a limited amount of ordering structure.

When systemd starts up the system, it will activate all units that are dependencies of default.target (as well as recursively all dependencies of these dependencies). Usually, default.target is simply an alias of graphical.target or multi-user.target, depending on whether the system is configured for a graphical UI or only for a text console. To enforce minimal ordering between the units pulled in, a number of well-known target units are available, as listed on **systemd.special(7)**.

The following chart is a structural overview of these well-known units and their position in the boot-up logic. The arrows describe which units are pulled in and ordered before which other units. Units near the top are started before units nearer to the bottom of the chart.



```

v v | v rescue.target
timers.target paths.target | sockets.target
| | |
\_____ | _____ | _____ /
\|
v
basic.target
|
|_____ /| emergency.service
/| | |
| | | v
v v v emergency.target
display- (various system (various system
manager.service services services)
| required for |
| graphical UIs) v
| | multi-user.target
| |
\_____ | _____ /
\|
v
graphical.target

```

Target units that are commonly used as boot targets are *emphasized*. These units are good choices as goal targets, for example by passing them to the *systemd.unit=* kernel command line option (see [systemd\(1\)](#)) or by symlinking *default.target* to them.

BOOTUP IN THE INITIAL RAM DISK (INITRD)


The initial RAM disk implementation (*initrd*) can be set up using *systemd* as well. In this case, boot up inside the *initrd* follows the following structure.

The default target in the *initrd* is *initrd.target*. The bootup process begins identical to the system manager bootup (see above) until it reaches *basic.target*. From there, *systemd* approaches the special target *initrd.target*. If the root device can be mounted at */sysroot*, the *sysroot.mount* unit becomes active and *initrd-root-fs.target* is reached. The service *initrd-parse-etc.service* scans */sysroot/etc/fstab* for a possible */usr* mount point and additional entries marked with the *x-initrd.mount* option. All entries found are mounted below */sysroot*, and *initrd-fs.target* is reached. The service *initrd-cleanup.service* isolates to the *initrd-switch-root.target*, where cleanup services can run. As the very last step, the *initrd-switch-root.service* is activated, which will cause the system to switch its root to */sysroot*.

```

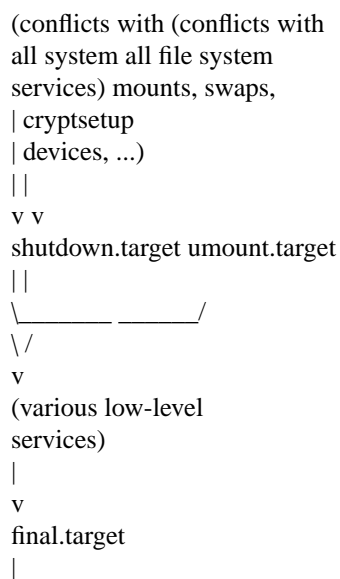
: (beginning identical to above)
:
v
basic.target
| emergency.service
|_____ /|
/| v
| sysroot.mount emergency.target
| |
| v
| initrd-root-fs.target
| |
| v
v initrd-parse-etc.service
(custom initrd |
services...) v
| (sysroot-usr.mount and

```



SYSTEM MANAGER SHUTDOWN

System shutdown with systemd also consists of various target units with some minimal ordering structure applied:



```
/||\  
|||\  
v v v v  
systemd-reboot.service systemd-poweroff.service systemd-halt.service systemd-kexec.service  
|||\  
v v v v  
reboot.target poweroff.target halt.target kexec.target
```

Commonly used system shutdown targets are *emphasized*.

SEE ALSO

[systemd\(1\)](#), [boot\(7\)](#), [systemd.special\(7\)](#), [systemd.target\(5\)](#), [dracut\(8\)](#)