

Bootup and Initialization

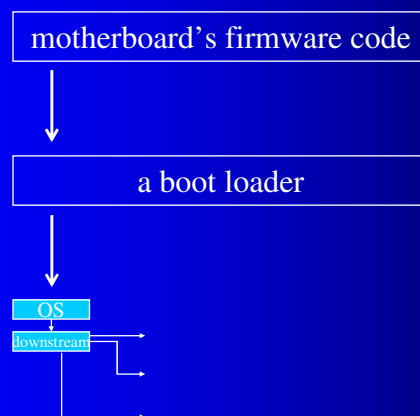
produced and directed by the Linux operating system

all about

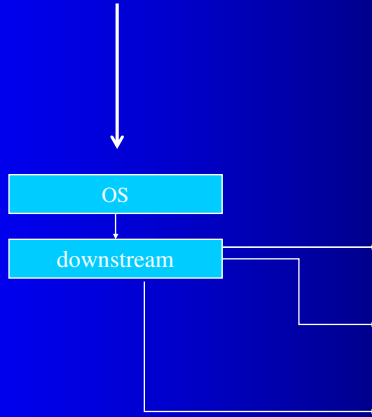
- SysV initialization method aka SysVinit (venerable, unix)
- systemd (newer, linux ~2010)

David Morgan

General PC Booting sequence



What happens "downstream"?



... it depends on the operating system

...and what if the OS is linux?

kernel v2.6.27 (fedora 10)

```
[root@frausto ~]# man boot | col -b | grep -B 1 -A 12 "kernel startup"
Kernel Startup
When the kernel is loaded, it initializes the devices (via their drivers), starts the swapper (it is a "kernel process", called kswampd in modern Linux kernels), and mounts the root file system (/).

Some of the parameters that may be passed to the kernel relate to these activities (e.g.: You can override the default root file system). For further information on Linux kernel parameters read bootparam(7).

Only then the kernel creates the first (user land) process which is numbered 1. This process executes the program /sbin/init, passing any parameters that weren't handled by the kernel already.

[root@frausto ~]# ls -l /sbin/init
lrwxrwxr-x 2 root root 138628 2008-05-13 08:27 /sbin/init
[root@frausto ~]# file /sbin/init
/sbin/init: ELF 32-bit LSB executable, Intel 80386, version 1 (SYSV), dynamically linked (uses shared libs), For GNU/Linux 2.6.9, stripped
[root@frausto ~]# man init | col -b | head -n 10
init(8)

NAME
    init - process management daemon

SYNOPSIS
    init [OPTION]...
```

/sbin/init gets executed hardcoded in the kernel

kernel v3.11.10 (fedora 20)

```
[root@unexgate ~]# man boot | grep -B 1 -A 12 "kernel startup"
kernel startup
When the kernel is loaded, it initializes the devices (via their drivers), starts the swapper (it is a "kernel process", called kswampd in modern Linux kernels), and mounts the root file system (/).

Some of the parameters that may be passed to the kernel relate to these activities (e.g.: You can override the default root file system). For further information on Linux kernel parameters read bootparam(7).

Only then the kernel creates the first (user land) process which is numbered 1. This process executes the program /sbin/init, passing any parameters that weren't handled by the kernel already.

[root@unexgate ~]# ls -l /sbin/init
lrwxrwxr-x 1 root root 22 Feb 10 2014 /sbin/init -> ../lib/systemd/systemd
[root@unexgate ~]# file /sbin/init
/sbin/init: symbolic link to ../lib/systemd/systemd
[root@unexgate ~]# man init | head -n 12
SYSTEMD(1)

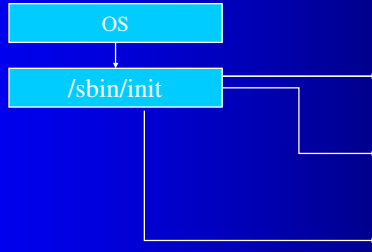
NAME
    systemd, init - system system and service manager

SYNOPSIS
    systemd [OPTIONS...]

    init [OPTIONS...] [COMMAND]

[root@unexgate ~]#
```

so with linux,
"downstream" == /sbin/init



and what does /sbin/init do?

kernel v2.6.27 (fedora 10)

```
[root@frausto ~]# man boot | col -b | grep -B 1 -A 12 "kernel startup"
kernel startup
When the kernel is loaded, it initializes the devices (via their drivers), starts the swapper (it is a "kernel process", called kswamp in modern Linux kernels), and mounts the root file system (/).
Some of the parameters that may be passed to the kernel relate to these activities. (e.g.: You can override the default root file system). For further information on Linux kernel parameters read bootparam(7).
Only then the kernel creates the first (user land) process which is numbered 1. This process executes the program /sbin/init, passing any parameters that weren't handled by the kernel already.

[root@frausto ~]# ls -l /sbin/init
-rwxr-xr-x 1 root root 128628 2008-05-13 08:27 /sbin/init
[root@frausto ~]# file /sbin/init
/sbin/init: ELF 32-bit LSB executable, Intel 80386, version 1 (SYSV), dynamically linked (uses shared libs), For GNU/Linux 2.6.9, stripped
[root@frausto ~]# man init | col -b | head -n 10
init(8)

NAME
  init - process management daemon

SYNOPSIS
  init [OPTION]...
```

kernel v3.11.10 (fedora 20)

```
[root@unexgate ~]# man boot | grep -B 1 -A 12 "kernel startup"
kernel startup
When the kernel is loaded, it initializes the devices (via their drivers), starts the swapper (it is a "kernel process", called kswamp in modern Linux kernels), and mounts the root file system (/).
Some of the parameters that may be passed to the kernel relate to these activities (e.g.: You can override the default root file system). For further information on Linux kernel parameters read bootparam(7).
Only then the kernel creates the first (user land) process which is numbered 1. This process executes the program /sbin/init, passing any parameters that weren't handled by the kernel already.

[root@unexgate ~]# ls -l /sbin/init
-rwxr-xr-x 1 root root 22 Feb 10 2014 /sbin/init -> ../lib/systemd/systemd
[root@unexgate ~]# file /sbin/init
/sbin/init: symbolic link to ../lib/systemd/systemd
[root@unexgate ~]# man init | head -n 12
SYSTEMD(1)
SYSTEMD(1)

NAME
  systemd, init - systemd system and service manager

SYNOPSIS
  systemd [OPTIONS...]

init [OPTIONS...] [COMMAND]
```

Whatever it's written to do usually run some setup stuff then call a shell

these /sbin/init's differ!

Difference: sysVinit vs systemd init

kernel v2.6.27 (fedora 10)

```

[roothfrausto ~]# man boot | col -b | grep -B 1 -A 12 "kernel startup"
Kernel Startup
When the kernel is loaded, it initializes the devices (via their drivers), starts the swapper (it is a "kernel process", called kswapper in modern Linux kernels), and mounts the root file system (/).

Some of the parameters that may be passed to the kernel relate to these activities (e.g.: You can override the default root file system). For further information on Linux kernel parameters read bootparam(7).

Only then the kernel creates the first (user land) process which is numbered 1. This process executes the program /sbin/init, passing any parameters that weren't handled by the kernel already.

[roothfrausto ~]# ls -l /sbin/init
lrwxrwxr-x 1 root root 138628 2008-05-13 08:27 /sbin/init
[roothfrausto ~]# file /sbin/init
/sbin/init: ELF 32-bit LSB executable, Intel 80386, version 1 (SYSV), dynamically linked (uses shared libs), For GNU/Linux 2.6.9, stripped
[roothfrausto ~]# man init | col -b | head -n 10
init(8)

NAME
    init - process management daemon

SYNOPSIS
    init [OPTIONS]...

[roothfrausto ~]#
    
```

an /sbin/init that performs sysVinit

an /sbin/init that does systemd init

kernel v3.11.10 (fedora 20)

```

[roothunexgate ~]# man boot | grep -B 1 -A 12 "kernel startup"
kernel startup
When the kernel is loaded, it initializes the devices (via their drivers), starts the swapper (it is a "kernel process", called kswapper in modern Linux kernels), and mounts the root file system (/).

Some of the parameters that may be passed to the kernel relate to these activities (e.g.: You can override the default root file system). For further information on Linux kernel parameters read bootparam(7).

Only then the kernel creates the first (user land) process which is numbered 1. This process executes the program /sbin/init, passing any parameters that weren't handled by the kernel already.

[roothunexgate ~]# ls -l /sbin/init
lrwxrwxr-x 1 root root 22 Feb 10 2014 /sbin/init -> ../lib/systemd/systemd
[roothunexgate ~]# file /sbin/init
/sbin/init: symbolic link to ../lib/systemd/systemd'
[roothunexgate ~]# man init | head -n 12
SYSTEMD(1)

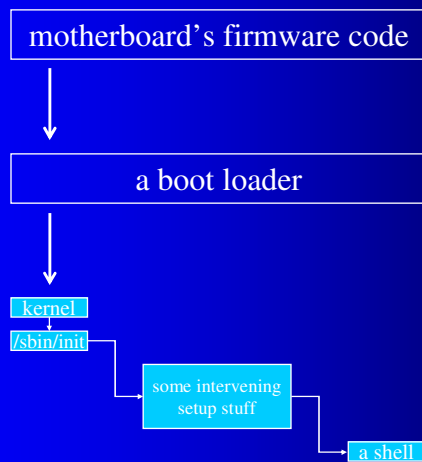
NAME
    systemd, init - systemd system and service manager

SYNOPSIS
    systemd [OPTIONS...]

    init [OPTIONS...] [COMMAND]

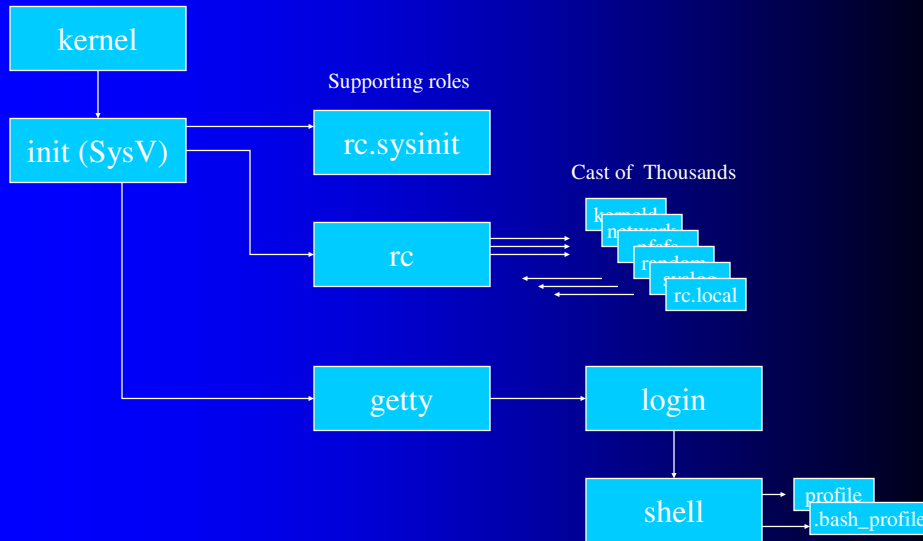
[roothunexgate ~]#
    
```

linux booting sequence

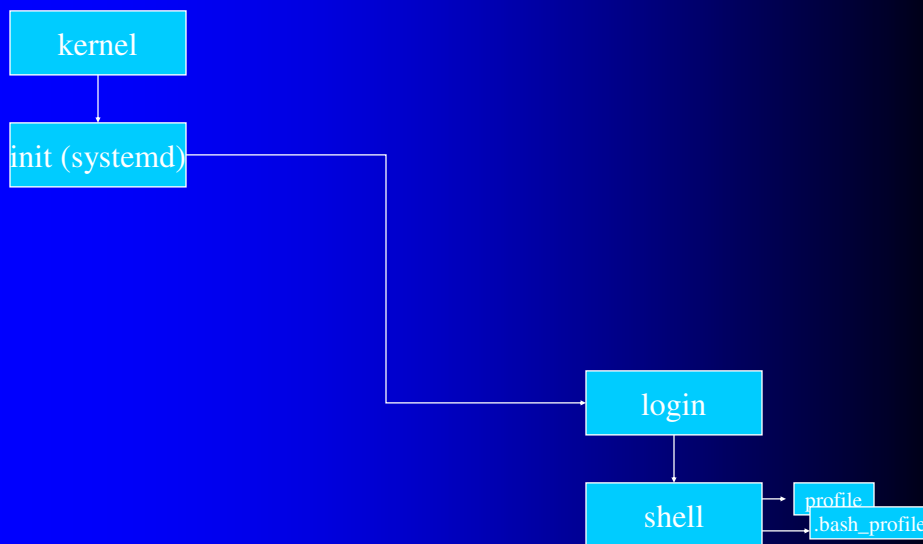


A Play in Several Acts - sysvinit

Starring roles



A Play in Several Acts - systemd



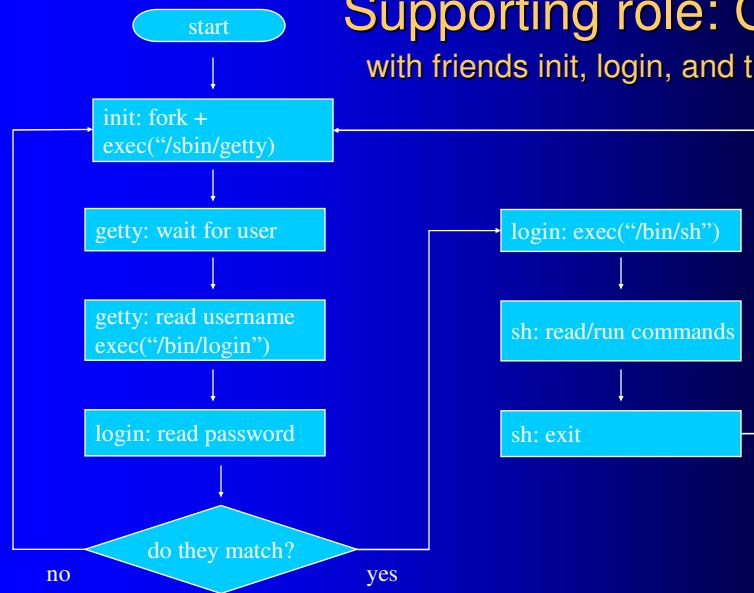
Starring role: Kernel

- kernel code loads (from somewhere)
 - from storage device with no filesystem
 - starting from first sector
 - from storage device with filesystem (most common)
 - from a file, e.g., /boot/vmlinuz
 - from a network
 - with PXE boot ("Preboot eXecution Environment")
- identifies/initializes hardware
- displays/stores messages ("dmesg" command shows)
- invokes init process

Starring role: init Process

- father of all processes
 - init is to process structure as root is to file structure
 - always PID number "1"
- creates other processes

Supporting role: Getty with friends init, login, and the shell



Shell Startup Files

- executed by shell when started by login
- `/etc/profile`, runs 1st
 - universal settings, all users
- `/home/username/.bash_profile`, runs 2nd
 - settings specific to user “*username*”

Turn services on/off manually (SysVinit)

- starting
 - /etc/rc.d/init.d/*<script for service>* start *or*
 - **service <script for service> start**
- stopping
 - /etc/rc.d/init.d/*<script for service>* stop *or*
 - **service <script for service> stop**
- services re-read configuration files when restarted
(restart one any time you change its config file)

Set services to boottime auto-on/off (SysVinit)

- set it to turn on
 - chkconfig <script for service> on**
- set it to not turn on
 - chkconfig <script for service> off**

So much for SysVinit.

Now what about systemd?

systemd

- open source project by Lennart Pottering
- relationship to SysV init system
 - a drop-in replacement (and more)
 - coexists with sysvinit in Fedora 15 (hybrid/transitional)
 - http://fedoraproject.org/wiki/SysVinit_to_Systemd_Cheatsheet
- parallelizes the numerous boot activities
- documentation
 - author's blog: <http://0pointer.de/blog>
 - see “14 May 2011” entry for links

sysvinit replacement role

Process Identifier 1

On every Unix system there is one process with the special process identifier 1. It is started by the kernel before all other processes and is the parent process for all those other processes that have nobody else to be child of. Due to that it can do a lot of stuff that other processes cannot do. And it is also responsible for some things that other processes are not responsible for, such as bringing up and maintaining userspace during boot.

Historically on Linux the software acting as PID 1 was the venerable sysvinit package, though it had been showing its age for quite a while. Many replacements have been suggested, only one of them really took off: Upstart, which has by now found its way into all major distributions.

As mentioned, the central responsibility of an init system is to bring up userspace. And a good init system does that fast. Unfortunately, the traditional SysV init system was not particularly fast.

For a fast and efficient boot-up two things are crucial:

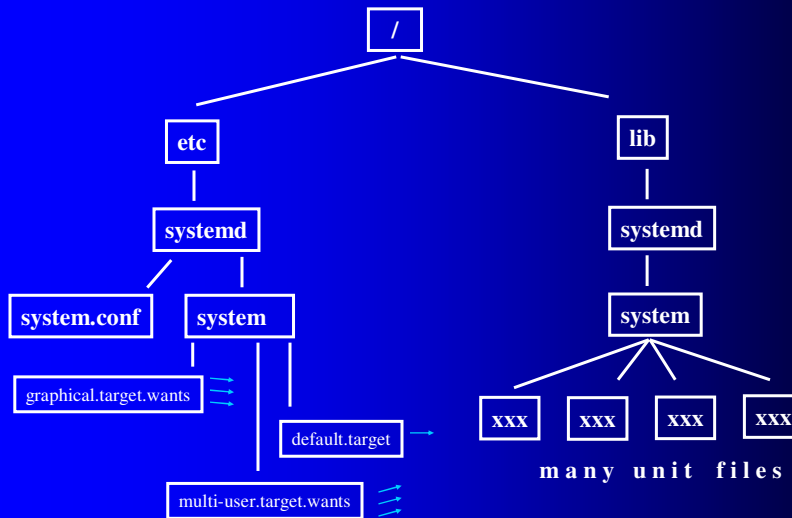
- To start **less**.
- And to start **more in parallel**.

What does that mean? Starting less means starting fewer services or deferring the starting of services until they are actually needed. There are some services where we know that they will be required sooner or later (syslog, D-Bus system bus, etc.), but for many others this isn't the case. For example, bluetoothd does not need to be running unless a bluetooth dongle is actually plugged in or an application wants to talk to its D-Bus interfaces. Same for a printing system: unless the machine physically is connected to a printer, or an application wants to print something, there is no need to run a printing daemon such as CUPS. Avahi: if the machine is not connected to a network, there is no need to run Avahi, unless some application wants to use its APIs. And even SSH: as long as nobody wants to contact your machine there is no need to run it, as long as it is then started on the first connection. (And admit it, on most machines where sshd might be listening somebody connects to it only every other month or so.)

Starting more in parallel means that if we have to run something, we should not serialize its start-up (as sysvinit does), but run it all at the same time, so that the available CPU and disk IO bandwidth is maxed out, and hence the overall start-up time minimized.

from Lennart Pottering <http://0pointer.de/blog/projects/systemd.html>

Default directories and files



a "default.target" unit file (as symlink)
and
a few "wants" subdirs containing symlinks

Default directories and files

```

root@localhost:~# ps ax | head -2
PID TTY          STAT TIME COMMAND
1 ?        Ss   0:02 /sbin/init
root@localhost:~# ls -l /sbin/init
-rwxr-xr-x 1 root root 16 Jul 11 17:45 /sbin/init -> ../lib/systemd
root@localhost:~# ls -l /etc/systemd/system
total 24
systemd.service          dev-hugepages.mount      graphical.target
accounts-daemon.service dev-hugepages.mount      halt-local.service
audit.service           dev-mqueue.mount         halt.service
blkid.service           display-manager.service  halt.target
brctl.service           emergency.service        httpd-daemon.target
brctl.target           fedora-autorelabel.mark.service
brctl.target.wants     fedora-autorelabel.service
bluetooth.service      fedora-autopop.service   irqbalance.service
bluetooth.target      fedora-configure.service kexec.service
bluetooth.target.wants fedora-configure.service kexec.target
brltty.service         fedora-landscape.service local-fs.target
brltty.target          fedora-remotely.service  local-fs.target.wants
brltty.target.wants   fedora-storage-init-late.service
brltty.service        fedora-storage-init-late.service
brltty.target          fedora-sysinit-unlock.service
brltty.target.wants   fedora-wait-storage.service
brltty.service        final.target
brltty.target          fireboot-graphical.service
brltty.target.wants  fireboot-graphical.service
brltty.service        firstboot-text.service
brltty.target          fire-root.service
brltty.service        flocks.service
brltty.target          getty.service
brltty.target.wants  getty.target
brltty.service        gum.service
brltty.target.wants  ntpd.service
root@localhost:~#

```

← PID #1 is /bin/systemd

← unit configuration files

Centralized “systemctl” utility

- inspect and control state of systemd
- not to be confused with “sysctl” utility !!

Turn services on/off manually

(set "current" state)

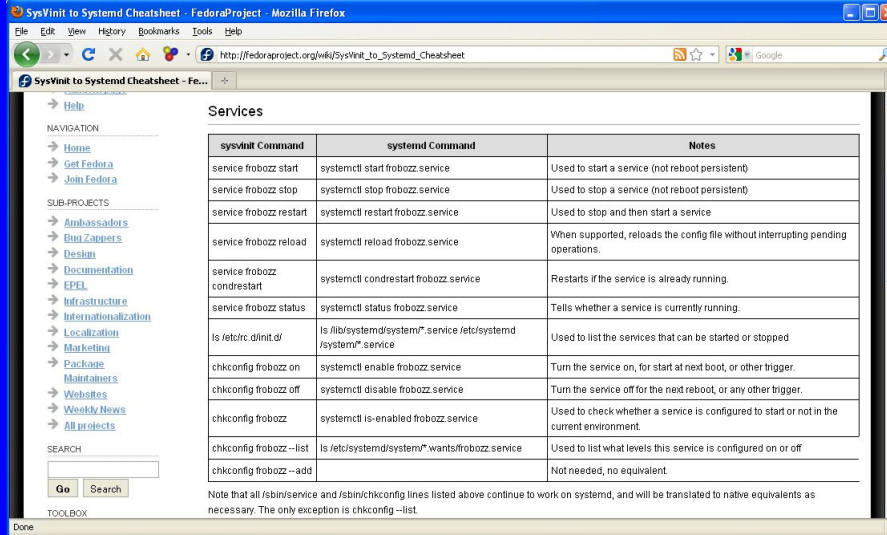
- starting
systemctl start *<service's unit file>*
- stopping
systemctl stop *<service's unit file>*
- example: turn on/off the logging service
systemctl start rsyslog.service
systemctl stop rsyslog.service

Set services to boottime auto-on/off

(set "persistent" state)

- set it to turn on
– **systemctl enable** *<service's unit file>*
- set it to not turn on
– **systemctl disable** *<service's unit file>*
- example: turn on/off the logging service
– systemctl enable rsyslog.service
– systemctl disable rsyslog.service

SysVinit to systemctl cheatsheet



The screenshot shows a web browser window displaying a cheatsheet for SysVinit to systemctl. The page is titled "SysVinit to Systemd Cheatsheet" and is part of the FedoraProject. The browser's address bar shows the URL "http://fedoraproject.org/wiki/SysVinit_to_Systemd_Cheatsheet". The page content includes a navigation menu on the left and a table of service management commands.

sysvinit Command	systemd Command	Notes
service frobozz start	systemctl start frobozz.service	Used to start a service (not reboot persistent)
service frobozz stop	systemctl stop frobozz.service	Used to stop a service (not reboot persistent)
service frobozz restart	systemctl restart frobozz.service	Used to stop and then start a service
service frobozz reload	systemctl reload frobozz.service	When supported, reloads the config file without interrupting pending operations.
service frobozz condrestart	systemctl condrestart frobozz.service	Restarts if the service is already running.
service frobozz status	systemctl status frobozz.service	Tells whether a service is currently running.
ls /etc/rc.d/init.d/	ls /lib/systemd/system*.service /etc/systemd/system*.service	Used to list the services that can be started or stopped
chkconfig frobozz on	systemctl enable frobozz.service	Turn the service on, for start at next boot, or other trigger.
chkconfig frobozz off	systemctl disable frobozz.service	Turn the service off for the next reboot, or any other trigger.
chkconfig frobozz	systemctl is-enabled frobozz.service	Used to check whether a service is configured to start or not in the current environment.
chkconfig frobozz --list	ls /etc/systemd/system* wants/frobozz.service	Used to list what levels this service is configured on or off
chkconfig frobozz --add		Not needed, no equivalent.

Note that all /sbin/service and /sbin/chkconfig lines listed above continue to work on systemd, and will be translated to native equivalents as necessary. The only exception is chkconfig --list.