

**Partitioning/FS/Mounting**

action	GPT
partition disk	gpart
init filesystem	newfs/mkfs
dev $\leadsto$ dir tree	mount

command	parameters
gpart	disk
newfs	partition, FS type
mount	partition, directory

**Partitioning (2)**

Should be done carefully (fixed sizes).

The system core should not be affected by file I/O of users.

$\leadsto$  /, /home, /var, /tmp should be on different file systems

swap at least as big as RAM

/var at least as big as RAM

**Partitioning (1)**

concept: additional layer between disk and FS

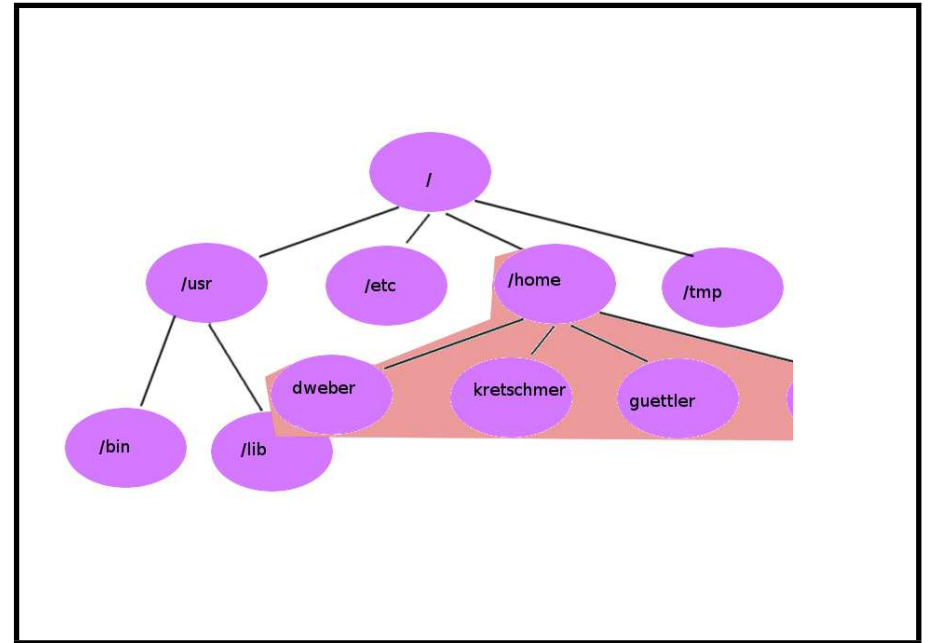
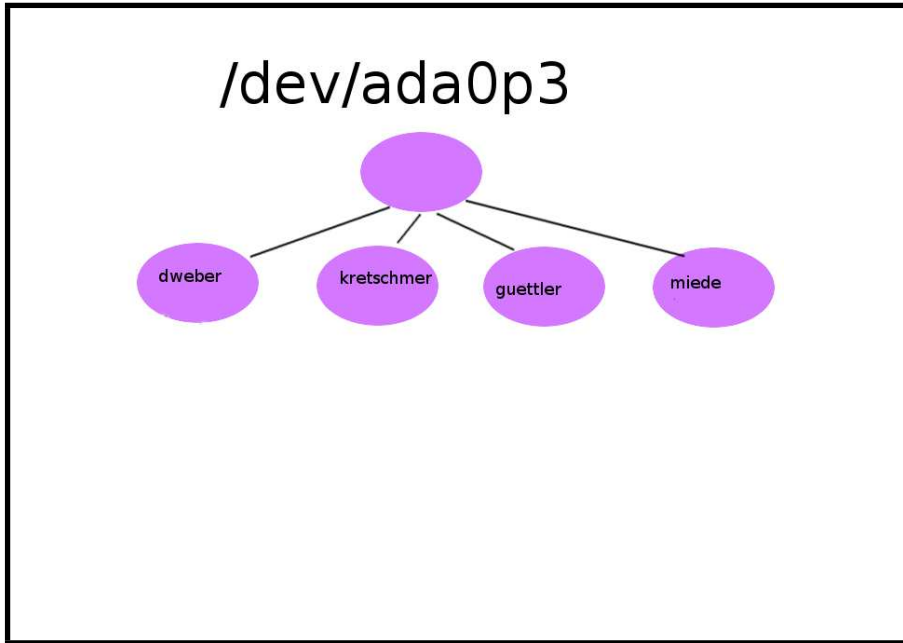


advantage:

- separated file storage
- controlled subsystems

disadvantage:

- fixed size (though growfs may resize)
- each partition to be configured

**Mounting an FS (1)**



### Mounting an FS (3)

Example:

```
# mount /dev/ada0p6 /tmp
```

Mounts partition /dev/ad0p6 as directory /tmp.

~>/tmp is called a mount point

~>mount point = empty directory

Mounting is usually done at boot time.

File /etc/fstab contains device-mount-mapping.

### Unmounting an FS (1)

Simple:

```
# umount /tmp
```

Or not so easy:

```
# umount /tmp
```

```
umount: unmount of /tmp failed: Device busy
```

We should *not* unmount an FS which is currently in use.

But we could:

```
# umount -f /tmp
```

This does *not* work for the *root filesystem*.

### /etc/fstab

# Device	M-point	FStype	Options	Dump	Pass#
/dev/ada0p2	/	ufs	rw	1	1
/dev/ada0p3	/usr	ufs	rw	2	2
/dev/ada0p4	/var	ufs	rw	2	2
/dev/ada0p5	/tmp	ufs	rw	2	2
/dev/ada0p10	/TMP	ufs	rw	2	2
134.96.216.92:/home	/home	nfs	rw	0	0
/dev/acd0	/cdrom	cd9660	ro,noauto	0	0

order of entriers important for mount, fsck

dump (# days), pass = order of FS check

### Unmounting an FS (2)

Which process uses a disk/file?

```
$ lsof | grep /home
COMMAND PID  USER  FD  TYPE      DEVICE SIZE/OFF  NODE NAME
bash    3627 dweber cwd  VDIR 255,117440514 1536 3379712 /home/dweber
lsof    3696 dweber cwd  VDIR 255,117440514 1536 3379712 /home/dweber
grep    3697 dweber cwd  VDIR 255,117440514 1536 3379712 /home/dweber
```

- alert corresponding users
- kill offending processes
- unmount the FS


## Filesystems on a RAMDISK

- create device node for this filesystem  
FreeBSD: mdconfig, OpenBSD/NetBSD: vnconfig,  
Solaris ramdiskadm
  - need info whether to use
    - \* simply allocated memory (malloc())
    - \* a file
    - \* swap space
  - need size
  - should provide a device number
- create filesystem on the device
- mount it

## Space Usage on file system: df = disk free

shows mounted file systems with

- name
- size in blocks (1K)
- number of used blocks
- number of available blocks
- percentage of use
- mount point



*Note:*

- must be checked periodically to avoid system failure
- likely overflows in /home, /var, /tmp

## Filesystems on a RAMDISK, Examples

using swap space

```
mdconfig -a -t swap -s 128M -u 10
newfs -U /dev/md10
mount /dev/md10 /tmp
chmod 1777 /tmp
```

using a file (with bsdlable)

```
dd if=/dev/zero of=somebackingfile bs=1k count=5k
mdconfig -a -t vnode -f somebackingfile -u 0
gpart create -s gpt md0
gpart add -t freebsd-ufs md0
newfs md0p1
mount /dev/md0p1 /mnt
```

## Example: df

Filesystem	1K-blocks	Used	Avail	Capacity	Mounted on
/dev/ada0p2	2063964	1068088	830760	56%	/
devfs	1	1	0	100%	/dev
/dev/ada0p4	4122780	1336824	2456136	35%	/var
/dev/ada0p5	4122780	208756	3584204	6%	/tmp
/dev/ada0p6	206417688	20924468	168979808	11%	/usr
/dev/ada0p7	20638108	9510384	9476676	50%	/home-local
isl1-03:/home	304689848	89136992	191177672	33%	/home
st1-ifs:/export/home_00	10737413240	724559112	10012859128	7%	/export/home_00

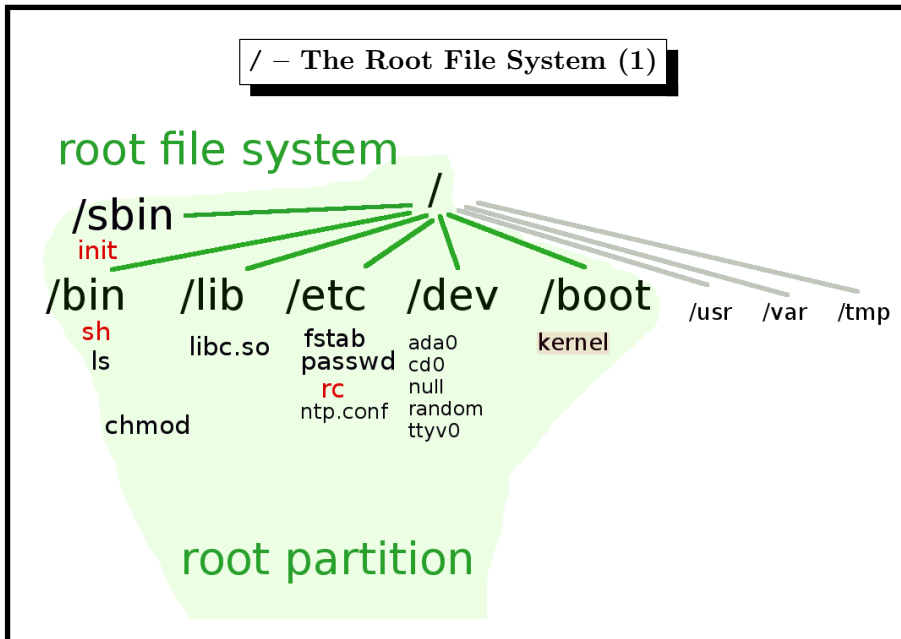
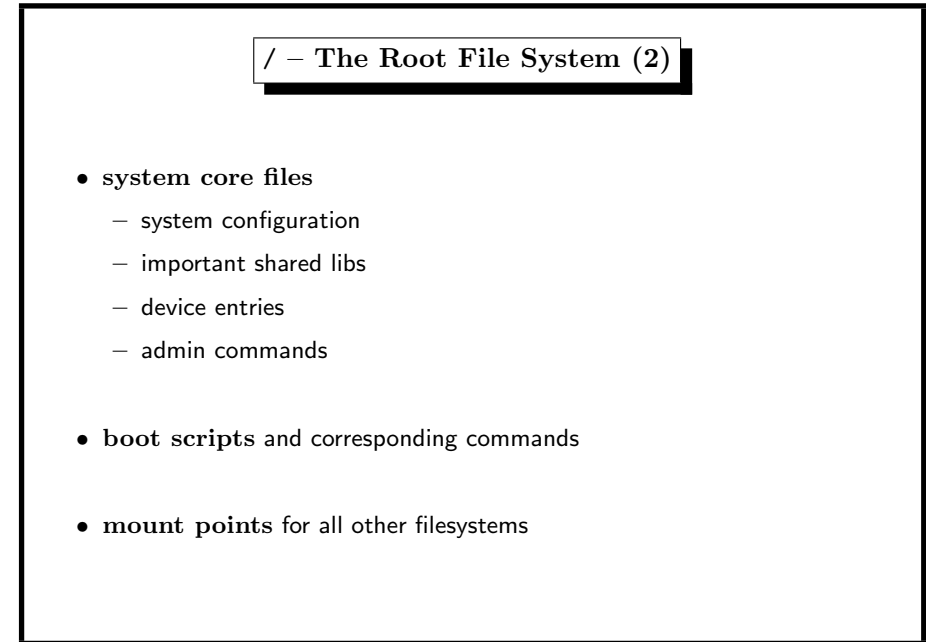
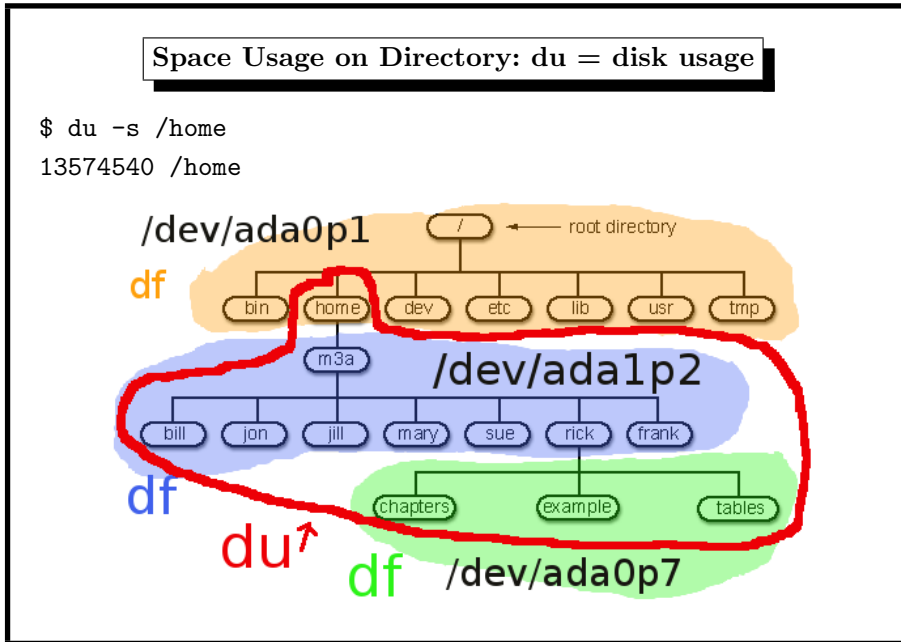
device

NFS-mounts

20G total

50% is in use

Mount  
point



### / – The Root File System (3)

Directory	Description	Example
<code>/bin</code>	user commands	<code>/bin/ls</code>
<code>/dev</code>	device entries	<code>/dev/ada0</code>
<code>/etc</code>	configuration	<code>/etc/passwd</code>
<code>/lib</code>	shared libraries	<code>/lib/libc.so</code>
<code>/sbin</code>	system administration commands	<code>/sbin/shutdown</code>
<code>/boot</code>	kernel binary, kernel modules	<code>/boot/kernel/kernel</code>
<code>(/proc)</code>	process information	<code>/proc/curproc/status</code>

## The Boot Problem

operating system does

- process management
- file system
- memory management
- I/O

but needs I/O and file system to read the operating system

- must determine system disk
- must read boot code from disk
- must read OS kernel from directory tree

→ chicken-and-egg problem

## Baron Münchhausen



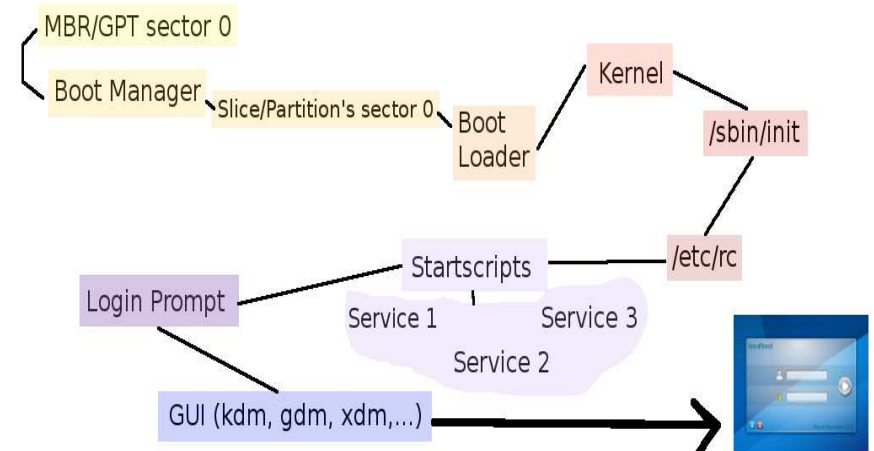
Es kann eben doch von Vorteil sein,  
wenn man einen gut trainierten Körper hat.

## Baron Münchhausen

Mein Pferd und ich wären hoffnungslos versunken,  
wenn ich es nicht geschafft hätte,  
mich an meinem eigenen Haarschopf aus dem Sumpf zu ziehen.



## Booting



## Starting to Boot (stage 0 boot)

- BIOS = basic input/output system, ROM...EEPROM...Flash
- BIOS locates MBR / GPT
- MBR/GPT code = *boot manager*, 512 bytes, boot menu
  - boot0, standard FreeBSD boot manager
  - GRUB,
  - standard PC MBR (searches active slice)
  - NTLDR, Vista MBR (Windows systems)
- MBR code reads boot loader (BIOS I/O)



## Boot Manager: Select Partition with a Root FS

FreeBSD boot0 start screen (file /boot/boot0, 512 bytes)

---


F1 DOS  
F2 FreeBSD  
F3 Linux  
F4 ??  
F5 Drive 1

Default: F2

---

source code directory /usr/src/sys/boot/i386/boot0


## Boot Manager / Boot Loader



a *boot manager* is independent from OS (on MBR)

a *boot loader* is OS specific (on slice)

## Prepare Loading of Boot Loader



FreeBSD boot1 (file /boot/boot1, 512 bytes)

Located in boot sector of bootable slice ~ 512 bytes.

Knows bsdlabel data structure.

Finds and loads boot2 (in the following 15 sectors)

## Locate Boot Loader on Partition

FreeBSD boot2 screenshot (file /boot/boot2, 7K bytes)

```
>> FreeBSD/i386 B00T
Default: 0:ad(0,a)/boot/loader
boot:
```


Knows how to find files on a UFS filesystem on it

until now, everything coded in machine language directly

Finds and loads /boot/loader, (217K)


## Side note: PXEBOOT

preboot-exec-environment (Intel), on ethernet card



~> diskless machines.

## Boot Loader: Prepare Loading of OS




/boot/loader

programmed in C, can do:

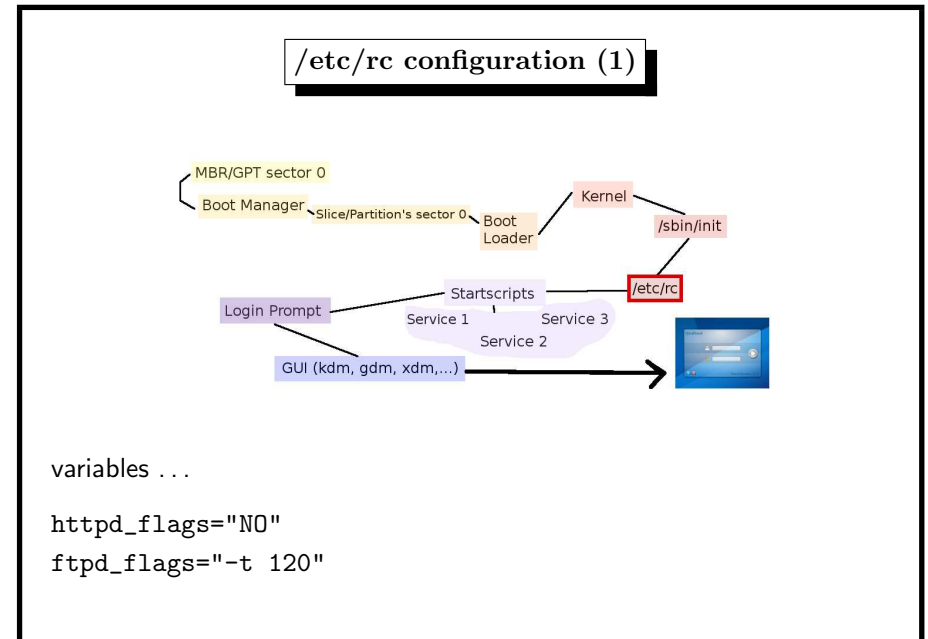
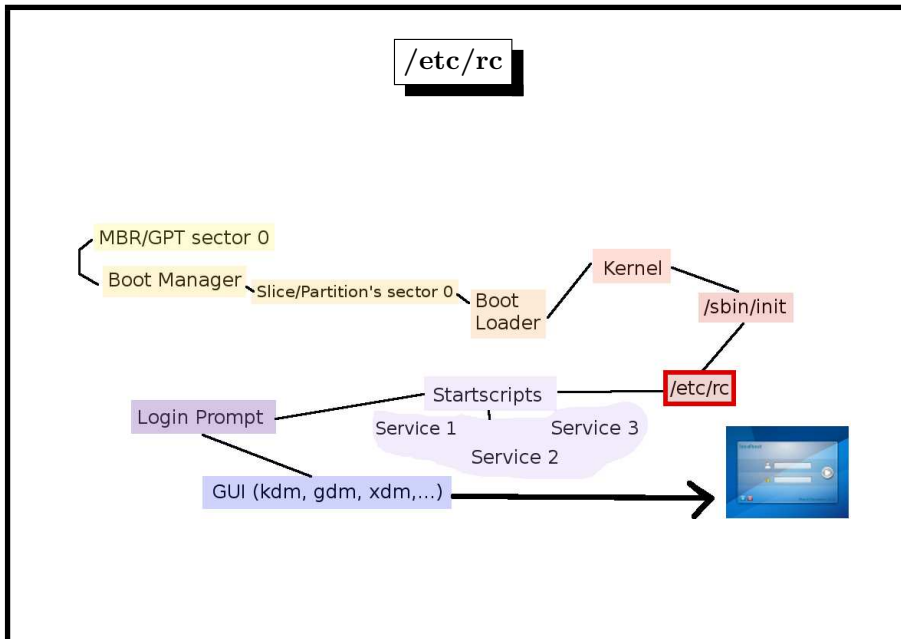
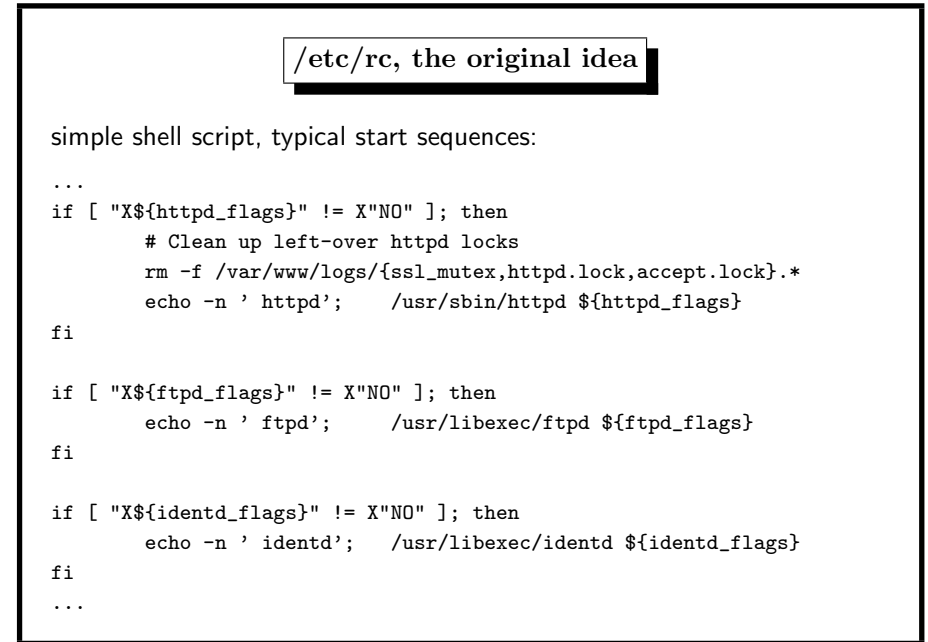
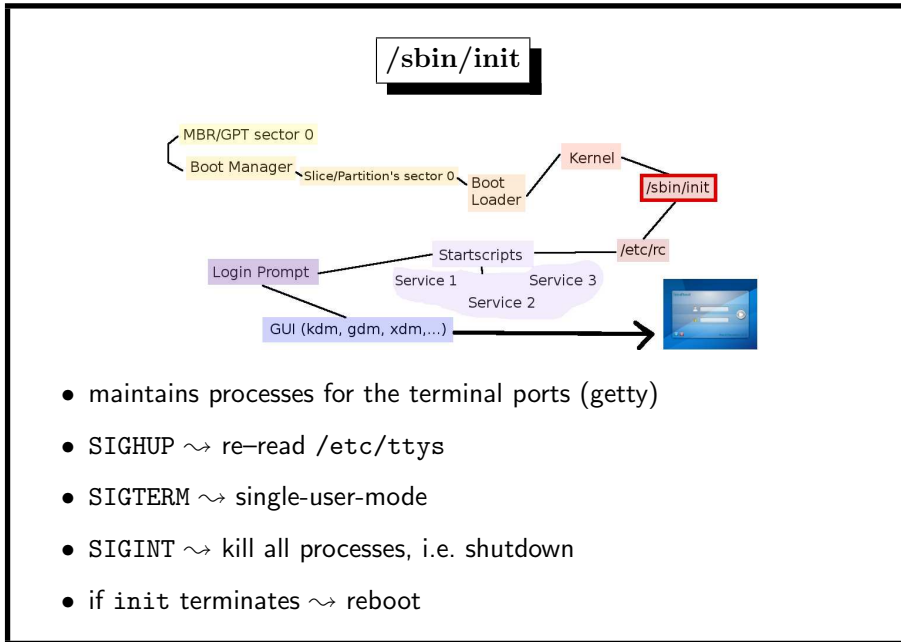
- probe for a console
- figure out what disk it is booting from
- probe for disks,
- load kernel/modules

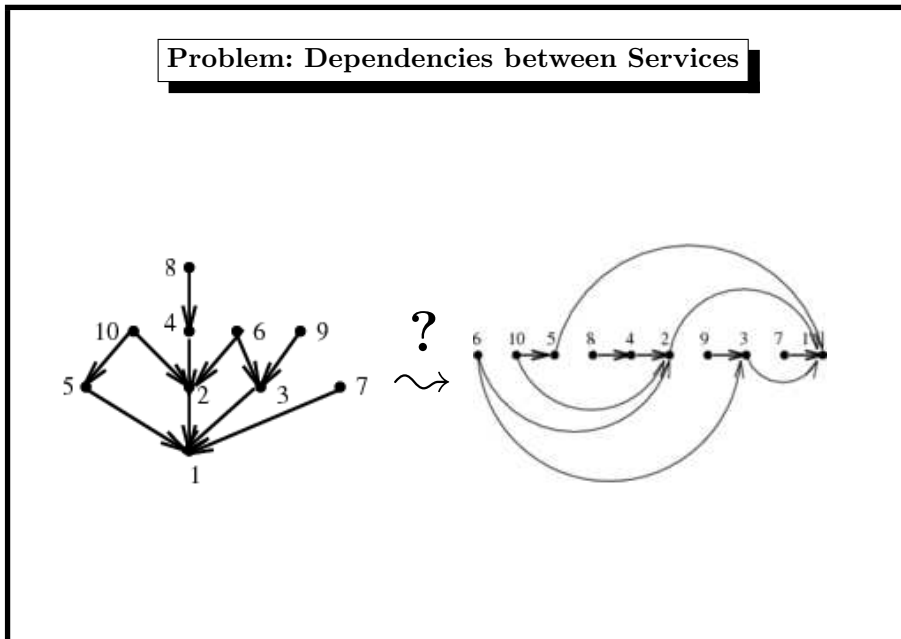
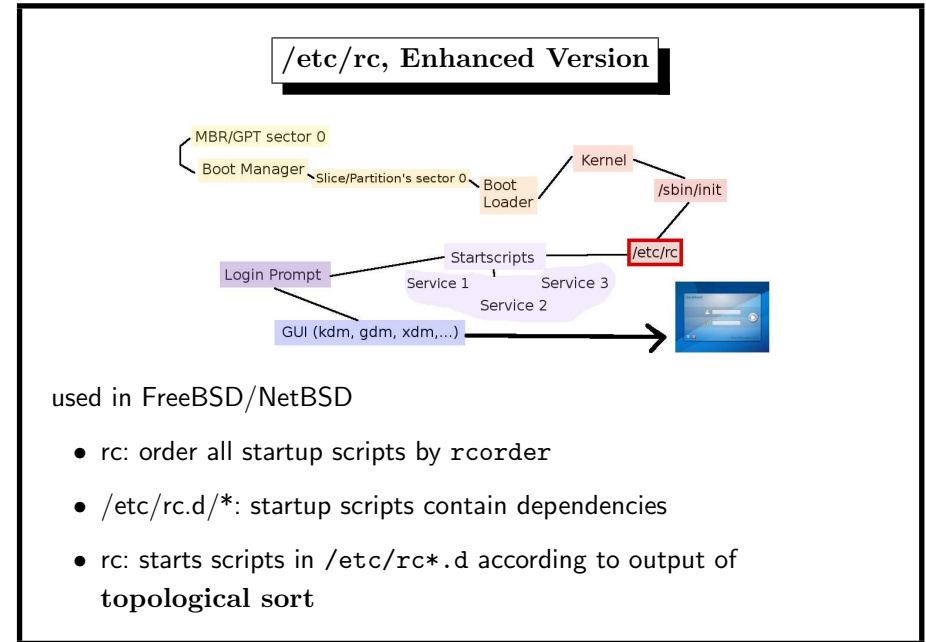
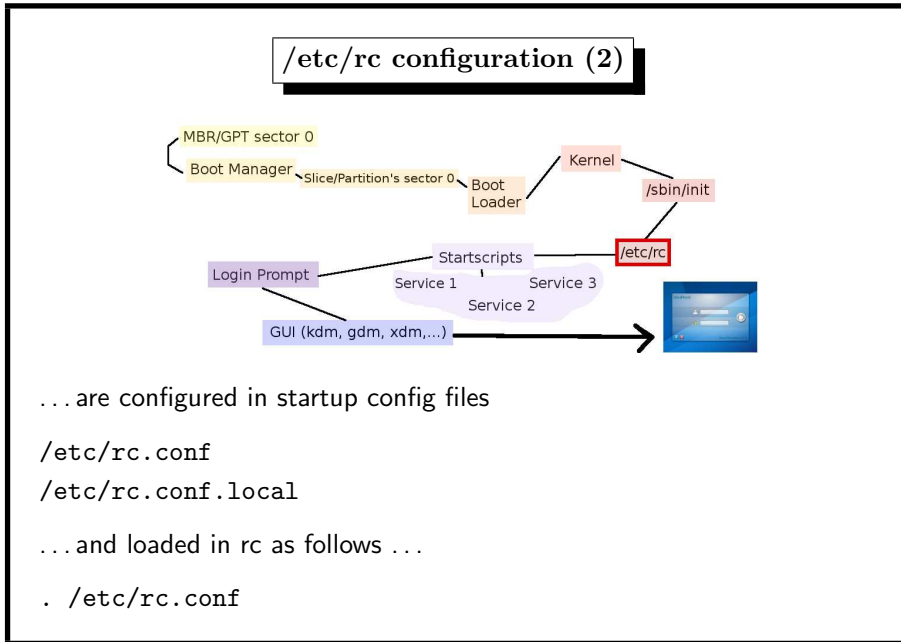
## Kernel



- initialize device drivers
- detect type of root filesystem, if unknown then STOP
- mount root filesystem read-only
- start process /sbin/init with PID=1
- init: /etc/rc, rc = resource configuration





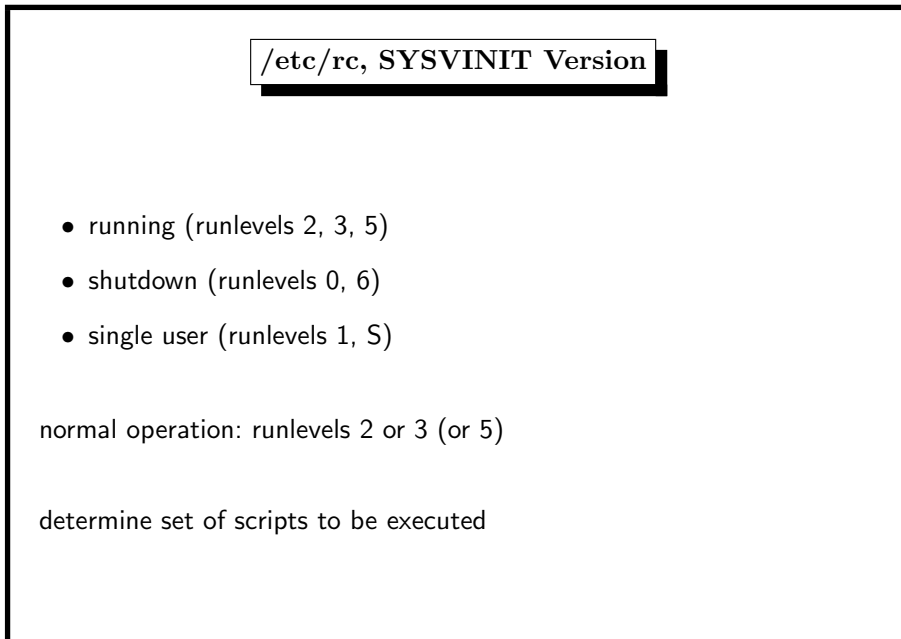
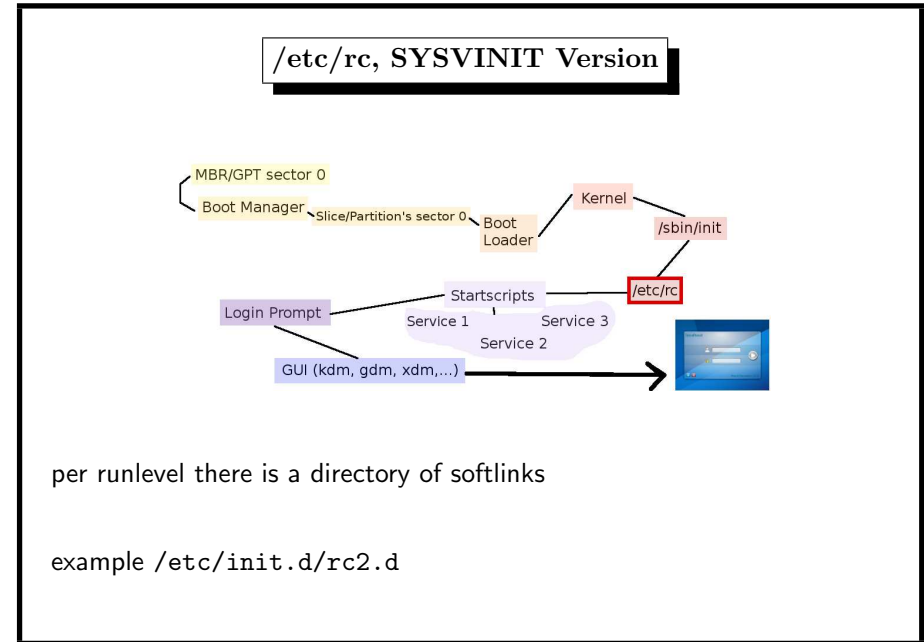
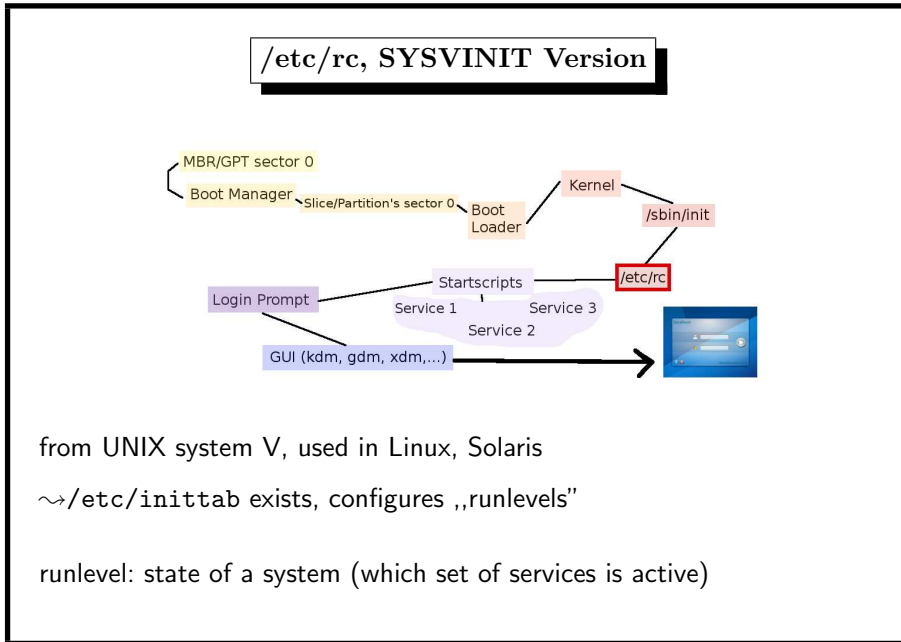


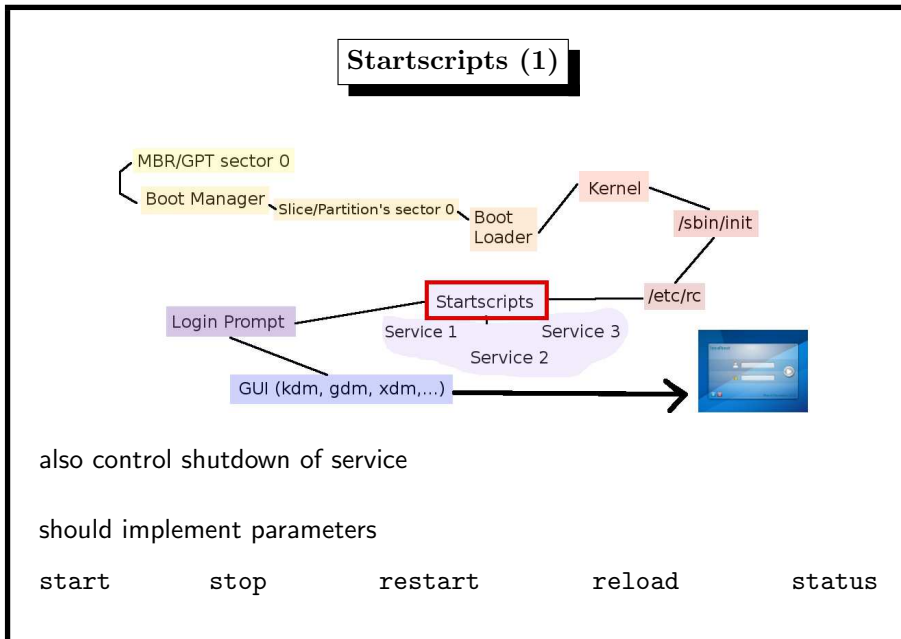
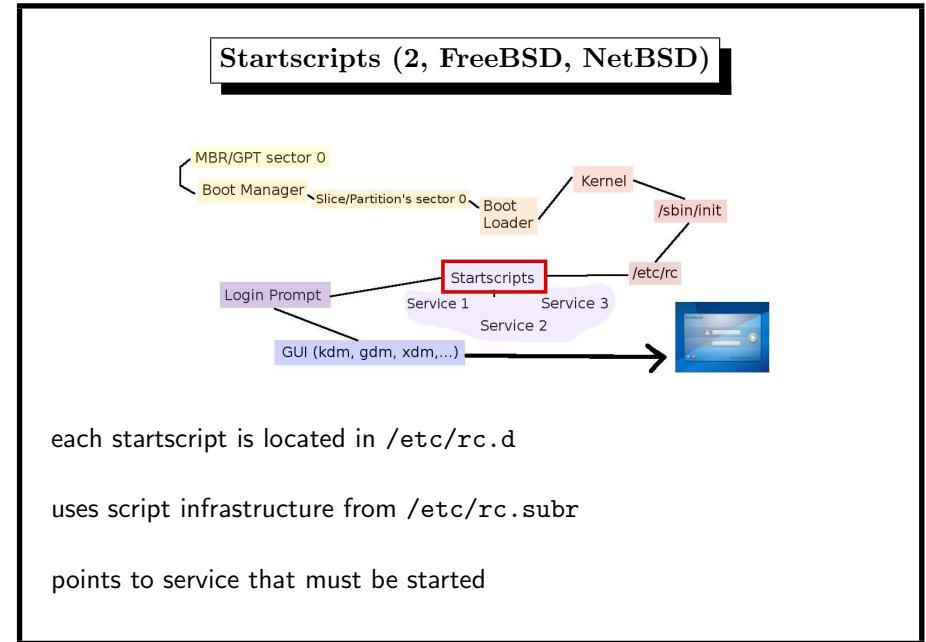
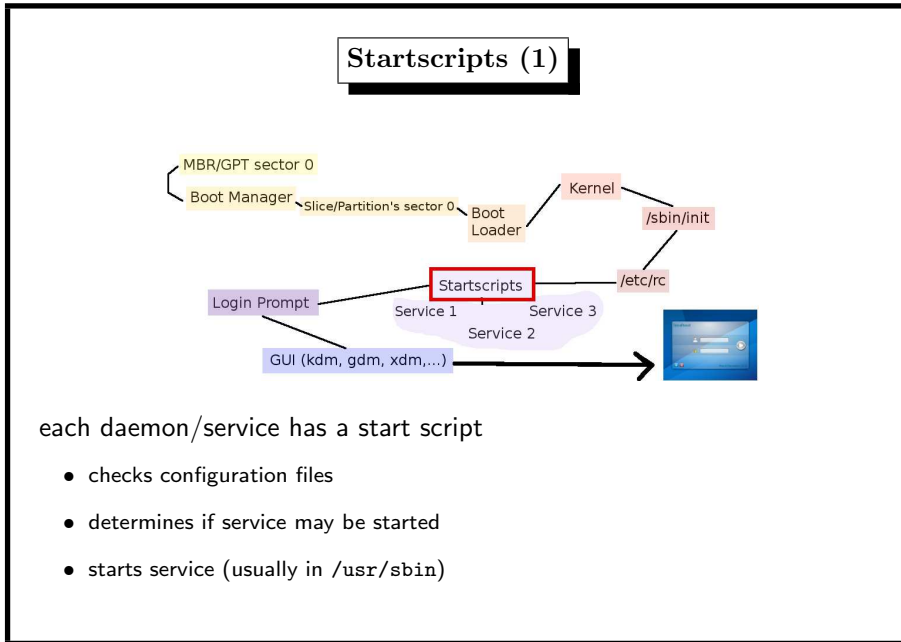
example: RPC service rpcbind

```
#!/bin/sh
#
```

```
# PROVIDE: rpcbind
```


```
# REQUIRE: NETWORKING ntpdate syslogd named
```





```
name="sshd"
rcvar='set_rcvar'
command="/usr/sbin/${name}"
start_precmd="sshd_precmd"
pidfile="/var/run/${name}.pid"
extra_commands="keygen reload"
```


## Single User Mode, Definition



- only root is allowed to log in
- only root filesystem is mounted


use this mode only for special tasks

## Invoking Single User Mode




- Use shutdown without `-h` or `-r`.
- On loader prompt use `boot -s`
- On loader menu use *single user*

## Single User Mode, Examples



- upgrade system (kernel, system lib, tools)
- repair filesystems after system crash
- forensics/clean-up after system break-in
- fix problems in critical system files
  - `/etc/fstab`
  - `/etc/inittab` (if SYSVINIT system)
- restore files from backup

## Login Prompt (text-oriented)




several text-oriented login screens

via `/etc/ttys`

```
# name  getty                type      status
ttyv0  "/usr/libexec/getty Pc"  cons2511  on  secure
ttyv1  "/usr/libexec/getty Pc"  cons2511  on  secure
ttyv2  "/usr/libexec/getty Pc"  cons2511  on  secure
...
```

- may be used to control root access to the machine (physical presence required)
- change resolution with
  - vidcontrol (FreeBSD)
  - (even 1024x768 resolution with MODE\_279)
  - kernel boot parameter (Linux)


### Side note: Console



a text-mode terminal (usually 80x25)

- usually used for root login
- direct connection to the motherboard
  - PS/2
  - serial
  - ...

### Side note: Console (2)




boot and have root ? FreeBSD-Version

see /etc/ttys on a FreeBSD-system


```
# If console is marked "insecure",
# then init will ask for the root password
# when going to single-user mode.
```

```
console none      unknown on insecure
```

### Login Prompt (for GUI)



### Side note: Console (3)



boot and have root ? Linux-Version

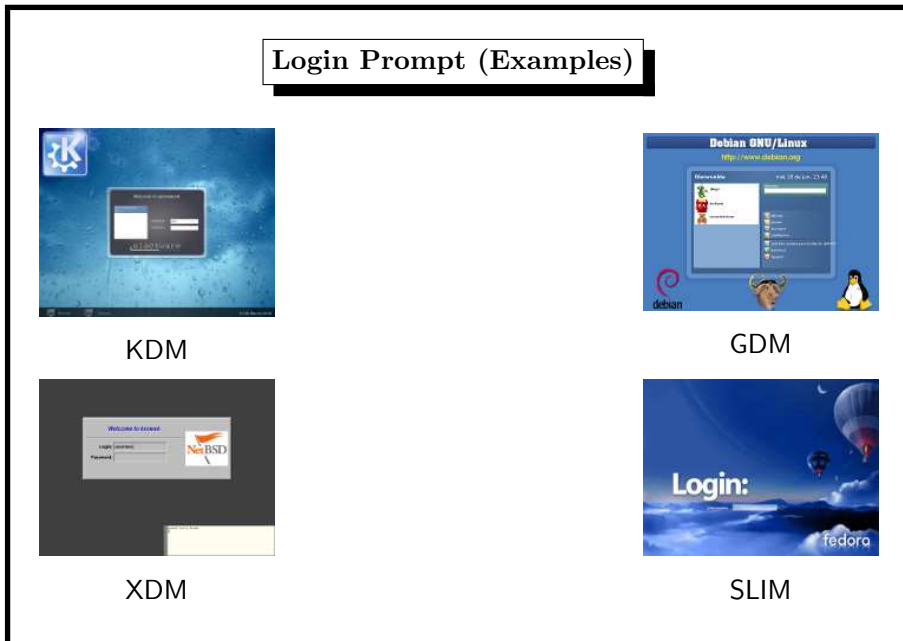

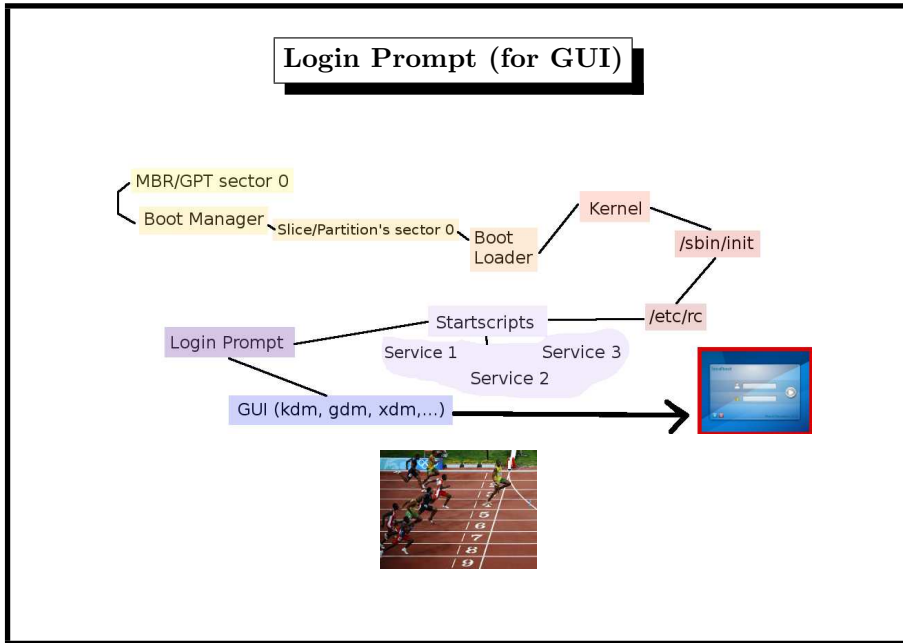
start from GRUB in single user mode

(append `single` on kernel-line and `init=/bin/bash`)

first process is root shell (no password needed)

~> must set password for GRUB/LILO

- depends on Xorg  
(GUI base system, formerly X11)
- requires root privileges (graphics card)
  - insecure: `SETUID /usr/local/bin/X` from terminal,
  - more secure: display manager  
(`xdm`, `kdm`, `gdm`, `slim`, ... as root)



### Load Average: How Busy the System Is

```

$ uptime
10:02AM up 31 days, 3:08, 3 users, load averages: 1.44 0.48 0.17
system time sessions
  
```

The screenshot shows the output of the `uptime` command. Annotations explain the load averages: `1.44` is the average number of processes ready to run over the last 15 minutes, `0.48` is the average over the last 5 minutes, and `0.17` is the average over the last minute. A yellow box contains the text: "uptime too small -> unstable server ?" and "uptime too big -> no security patches ?".