

LPIC-1 101-500 – Lesson 12

102.1 Design hard disk layout




Introduction

- During the installation of a new Linux system we are given the option to partition the hard disk to different filesystems
- Every filesystem occupies its own partition
- For personal computers, small servers or virtual machines we usually do not need more than 1 or 2 filesystems e.g. **/boot**, **/ (root)**
- For large systems like physical servers we generally have more filesystems




The *root* "/" filesystem

- It is the root of the filesystem tree and all other directories and filesystems are subdirectories of "/"
- It hosts the appropriate applications and files for the proper operation of a Linux system
- It usually does not need to be very large in case /var, /usr,
- /home and /tmp directories are hosted on their own partitions
- If the case above is true "/" is static in capacity and the minimum space required can be as low as **500MB** and not larger than **2GB**



The */boot* filesystem

- Hosted the kernel files (`vmlinuz*`), ramdisk files (`initrd*`), bootloader files (`/boot/grub`) and generally files related to booting up the system
 - It grows only in case new kernels are installed. In this case old kernels need to be removed so that `/boot` does not fill up
 - A typical size for `/boot` is up to **100MB** with maximum up to **500MB**
 - Old systems had a restriction where the BIOS could not read beyond 1024th cylinder. In that case `/boot` had to be equal or less than the size corresponding to the 1023th cylinder or else **BIOS** could not find a kernel to boot the system
- 

The */boot/efi* filesystem

- Known as the **EFI System Partition (ESP)**
- It is needed when the system has **UEFI** instead of a **BIOS** and the disks are **GPT** formatted instead of **MSDOS** (aka **MBR**)
- The minimum size for **/boot/efi** is **100MB** but **512MB** is recommended for compatibility with other systems (like Windows) on dual boot systems. The system is formatted as **FAT32**




The */home* filesystem

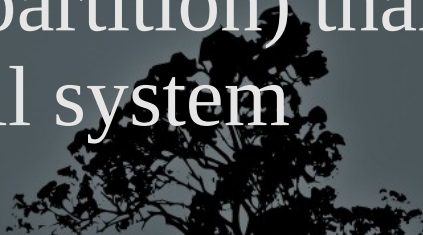
- Hosts the user home directories (homedirs)
- For users **user1** and **user2** the **/home/user1** and **/home/user2** directories are assigned respectively
- The environment variable **\$HOME** as well as the tilde "~" symbol, represent the home directory for the current user
- The size of **/home** depends on the number of users, the size of the user data and the available disk space on the disk. Generally it is one of the larger filesystems



The */usr* filesystem

- Hosts user (non-system) applications and initial data of applications, user libraries, as long as documentation for applications
(*/usr/share/doc/<application name>*)
 - It grows only when a new a new program, library or documentation is installed
 - In many setup this can be a common network directory (e.g. via NFS) for many systems
 - The size is determined from the size of the application files and the expected growth in case of new applications
- 

The */var* filesystem


- Hosts variable data like log files, emails, print spools, temporary application files, cached files, application data, databases, etc
 - It is one of the most rapidly changing filesystems
 - Its size depends on the specific needs. If there will be no data hungry applications it can be as small as **2GB**
 - It is usually recommended for the **/var** directory to be on a different filesystem (aka partition) than "/" so in case it fills up, the overall system stability is not affected
- 

The */tmp* filesystem

- Hosts temporary files for users, like session data, temporary application files, etc
- In this directory everybody has writing rights but can only delete the files/directories they own!
- If a user writes too many files it can fill the system up so it is recommended to exist as a different filesystem



The *swap* partition

- The **swap** partition has its own format
 - It serves as a temporary virtual memory in case the real memory (RAM) fills up
 - Traditionally there was a thumb rule that it should be double the RAM
 - In systems with large RAM (>2GB) it does not need to be bigger than 2GB but that depends on the system applications
 - Instead of a swap partition we can have a swap file hosted on a normal filesystem
 - Because swapping is related with degraded performance some systems with high I/O are build with big enough RAM and no swap
- 

Filesystem Mount Points

- Mount points are directories attached to filesystems
- They can be set on startup in the `/etc/fstab` file:
\$ cat /etc/fstab
- The active mount points can be found in `/etc/mtab`
\$ cat /etc/mtab
- **\$ mount #** show filesystems related to their mount points and their mount arguments
- **\$ df -hT #** shows filesystem usage, in human readable form (**-h**) and the type of the filesystem (**-T**)

Filesystem Mount Points

```
$ cat /etc/mtab
```

```
/dev/sda2 / ext3 rw 0 0
```

```
proc /proc proc rw 0 0
```

```
sysfs /sys sysfs rw 0 0
```

```
devpts /dev/pts devpts rw,gid=5,mode=620 0 0
```

```
/dev/sda8 /tmp ext3 rw 0 0
```

```
/dev/sda7 /var ext3 rw 0 0
```

```
/dev/sda6 /usr ext3 rw 0 0
```

```
/dev/sda5 /home ext3 rw 0 0
```


```
/dev/sda1 /boot ext3 rw 0 0
```

```
tmpfs /dev/shm tmpfs rw 0 0
```

```
none /proc/sys/fs/binfmt_misc binfmt_misc rw 0 0
```

```
sunrpc /var/lib/nfs/rpc_pipefs rpc_pipefs rw 0 0
```

Mount points are
marked in yellow



Partitions

- `$ cat /proc/partitions # show disks and their partitions`
- `# fdisk -l # show partitions, filesystems, their size and hard disk parameters`
- `# parted -l # similar to fdisk -l`

http://en.wikipedia.org/wiki/Master_boot_record
https://en.wikipedia.org/wiki/GUID_Partition_Table



Logical Volume Manager – LVM


- **LVM** provides more flexibility in managing storage
- **Physical Volumes (PV)**: the PV are the disks or partitions assigned to a Volume Group
- **Volume Group (VG)**: the VG is at the top of the LVM hierarchy. It receives storage resources in the form of PV and exports Logical Volumes to be used as filesystem
- **Logical Volumes (LV)**: logical volumes are logical structures, grouped under a VG and they replace the partitions of a more traditional storage

Logical Volume Manager – LVM

- LVM features:
- Logical volumes can grow online
- Volume groups can grow online by adding more PVs
- Supports snapshots
- Supports RAID (1, 5, 6) and striping
- High Availability features



License



The work titled "LPIC-1 101-500 – Lesson 12" by Theodotos Andreou is distributed with the Creative Commons Attribution ShareAlike 4.0 International License.

