

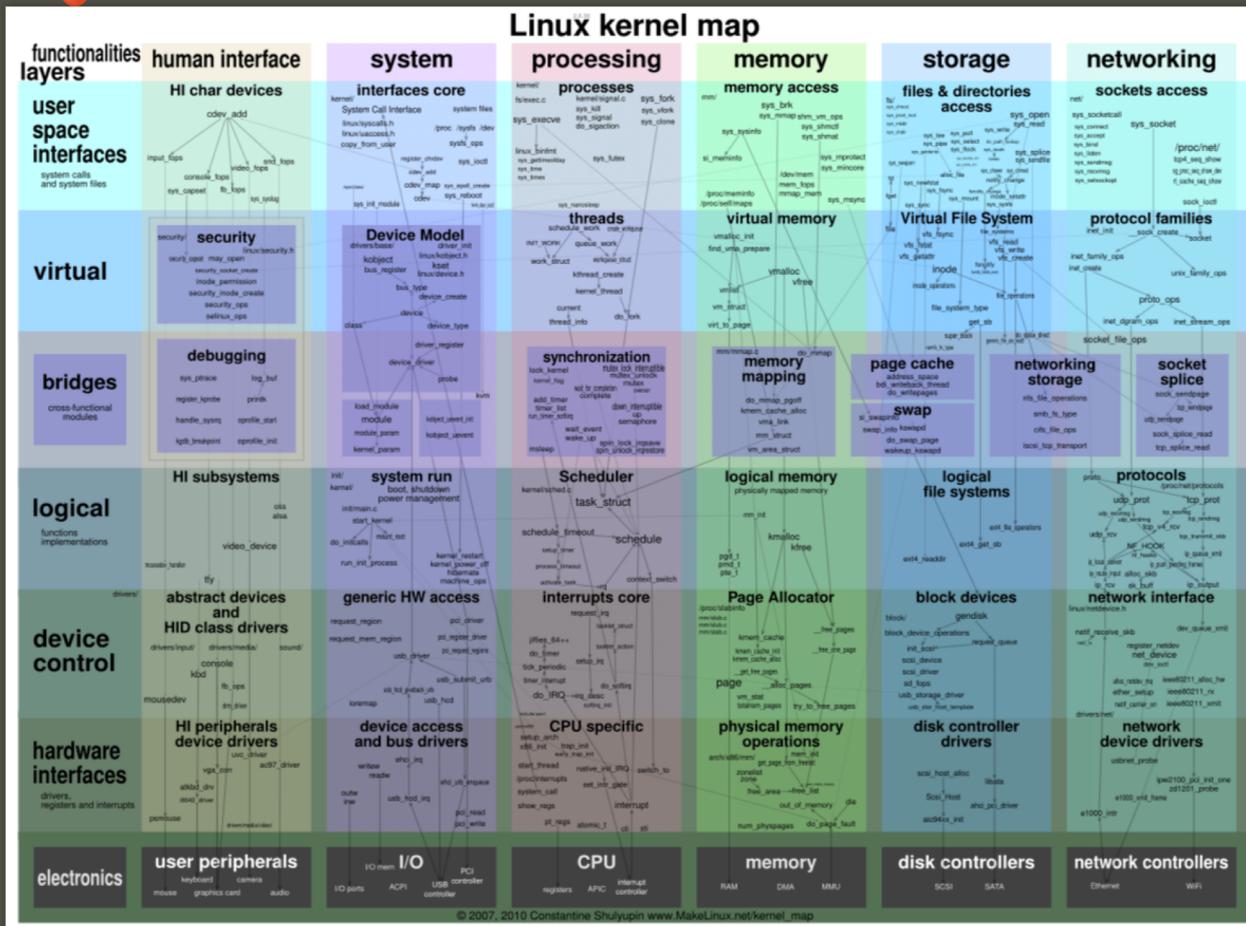
# Digging into Linux File Systems

# Not a goal of this presentation

- Rise interest in FS part of the Linux
- Make people more aware
- Have more fun with your favorite distro
- Hopefully learning something new

Trying to convince you to any specific FS!!

# Why interest in FS?



Source: [http://www.makelinux.net/kernel\\_map/](http://www.makelinux.net/kernel_map/)

# Standard Stack

EXT3, EXT4, XFS, ReiserFS

System Calls

VFS

File System

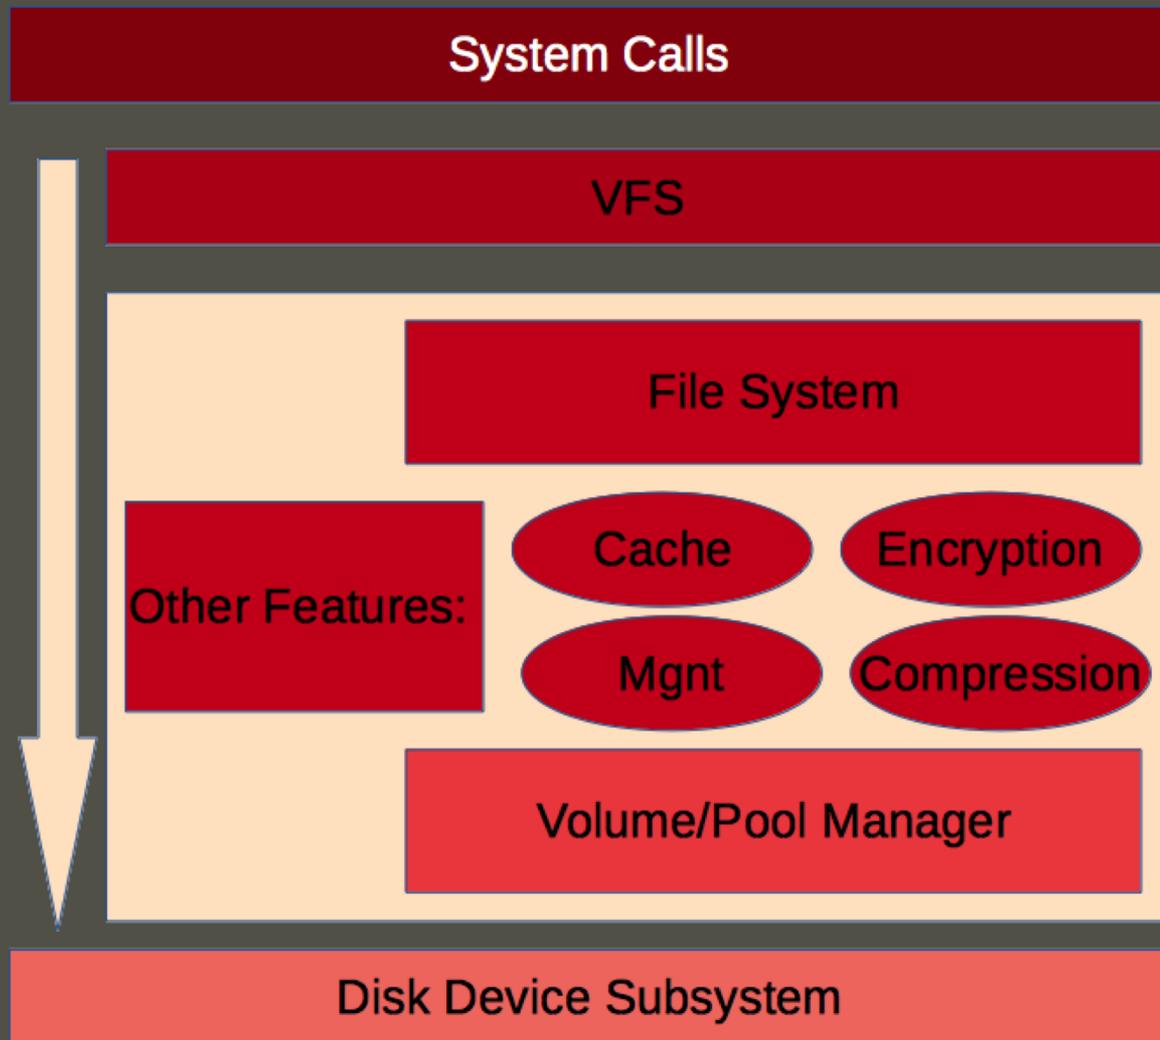
Volume Manager

Disk Device Subsystem



**More modern approach**

Examples: ZFS, BTRFS



# Difference in practice: separate VM

```
# LVM stuff
vgcreate lvm_vol /dev/sdb
lvcreate --name rootdg lvm_vol

#FS stuff
mkfs -t xfs /dev/mapper/lvm_vol-rootdg
mount -t xfs /dev/mapper/lvm_vol-rootdg /mnt

:/mnt# df -h
Filesystem                Size      Used Avail Use% Mounted on
/dev/mapper/lvm_vol-rootdg 3.7T      3.8G   3.7T   1% /mnt

:/# vgs
  VG          #PV #LV #SN Attr   VSize VFree
  lvm_vol     4   1   0 wz--n- 3.64t    0

:/# lvs
  LV          VG          Attr          LSize Pool Origin Data%  Meta%  Move Log
Cpy%Sync Convert
  rootdg     lvm_vol    -wi-ao----- 3.64t
```

# Difference in practice: VM inside FS

```
zpool create -f pool /dev/sdb
zfs set mountpoint=/mnt pool
zfs create pool/fs1
```

```
:-# zfs list
```

NAME	USED	AVAIL	REFER	MOUNTPOINT
pool	457K	3.84G	291K	/mnt
pool/fs1	19K	3.84G	19K	/mnt/fs1

```
:/mnt# df -T
```

Filesystem	Available	Use%	Mounted on	Type	1K-blocks	Used
pool				zfs	4031104	256
4030848	1%		/mnt			
pool/fs1				zfs	4030976	128
4030848	1%		/mnt/fs1			

# Managing Snapshots

- ZFS
- Create pool
- Create Data set (FS)

```
:/# cd /mnt/fs1
wget... - 'alice30.txt' saved [159332/159332]
:/mnt/fs1# zfs snapshot pool/fs1@testsnap1
:/mnt/fs1# zfs list -t snapshot
NAME                                USED  AVAIL  REFER  MOUNTPOINT
poolik/fs1@testsnap1                28K   -     116K   -

:/mnt/fs1# mv alice30.txt wonderland.txt
:/mnt/fs1# ls
wonderland.txt
:/mnt/fs1# echo "Hello Alice" >> wonderland.txt

:/mnt/fs1# mount -t zfs pool/fs1@testsnap1 /mnt1
:/mnt/fs1# ls /mnt1/
alice30.txt

:/mnt/fs1# mount | grep fs1
pool/fs1 on /mnt/fs1 type zfs (rw,relatime,xattr,noacl)
pool/fs1@testsnap1 on /mnt1 type zfs (ro,relatime,xattr,noacl)

:/mnt/fs1# diff /mnt1/alice30.txt ./wonderland.txt
3852a3853
> Hello Alice
```

# Managing Snapshots

## ○ LVM

```
/mnt# lvcreate -L 1M -s -n vol_snap /dev/mapper/lvm_vol-rootdg
```

```
Rounding up size to full physical extent 4.00 MiB
```

```
Volume group "lvm_vol" has insufficient free space  
(0 extents): 1 required.
```

```
# Fix the issue
```

```
lvdisplay lvm_vol
```

```
lvreduce -L 200G /dev/lvm_vol/rootdg
```

```
Resize2fs /dev/lvm_vol/rootdg 100G
```

```
lvcreate -L 1M -s -n vol_snap /dev/mapper/lvm_vol-rootdg
```

```
mount /dev/mapper/lvm_vol-vol_snap-cow /mnt1
```

## Compression

- Use ZFS  
(OpenZFS for Ubuntu)
- BTRFS also can be used here

```
zfs set compresssion=lz4 pool
```

```
zfs set compression=on pool
```

```
zfs get compressratio /mnt
```

NAME	PROPERTY	VALUE	SOURCE
poolik	compressratio	1.00x	-

```
wget alice-table.html
```

```
for i in `seq 1 1000`; do cat alice >> alice;  
done
```

```
ls -lahtri alice
```

```
13 -rw-r--r-- 1 u u 3.0M Apr 25 00:51 alice
```

```
zfs get used,logicalused poolik
```

NAME	PROPERTY	VALUE	SOURCE
poolik	used	435K	-
poolik	logicalused	3.40M	-

```
zfs get compressratio /mnt
```

NAME	PROPERTY	VALUE	SOURCE
poolik	compressratio	10.77x	-

# What is showing your mount command?

Thank you for attention