

RED HAT
SUMMIT

LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.

June 11-14, 2013
Boston, MA



RED HAT
SUMMIT

Fundamentals of LVM (LAB)

Christoph Doerbeck

Principal Solutions Architect, Red Hat

Jonathan Brassow

Principal Software Engineer, Red Hat

06.12.2013

RED HAT SUMMIT

Fundamentals of LVM

Please proceed to:

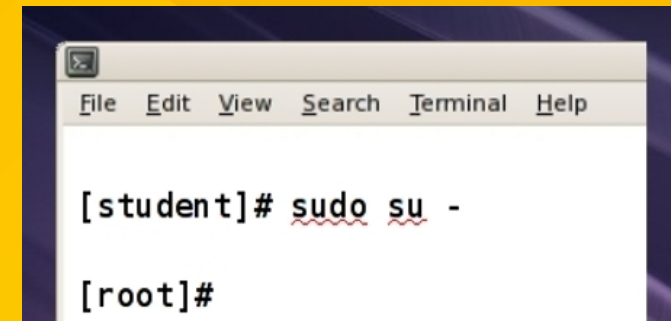

log in to workstation

log in to virtual machine

disable screen saver

disable screen lock

open a root shell



RED HAT
SUMMIT

Quick Survey

Audience Profile

System Administrators?

Programmers?

RHCSA's?

RHCE's?

RHCA's?

RED HAT SUMMIT

Quick Survey

BASH Skills Test

```
ls /dev/vd{a,b,c}
```

```
ls /dev/vda{1..3}
```

```
Ctrl-R
```

Quick Survey

BASH Skills Test

```
ls /dev/vd{a,b,c}
```

```
ls /dev/vda{1..3}
```

```
Ctrl-R
```

```
ls /dev/vda /dev/vdb /dev/vdc
```

```
ls /dev/vda1 /dev/vda2 /dev/vda3
```

```
Search command line history
```

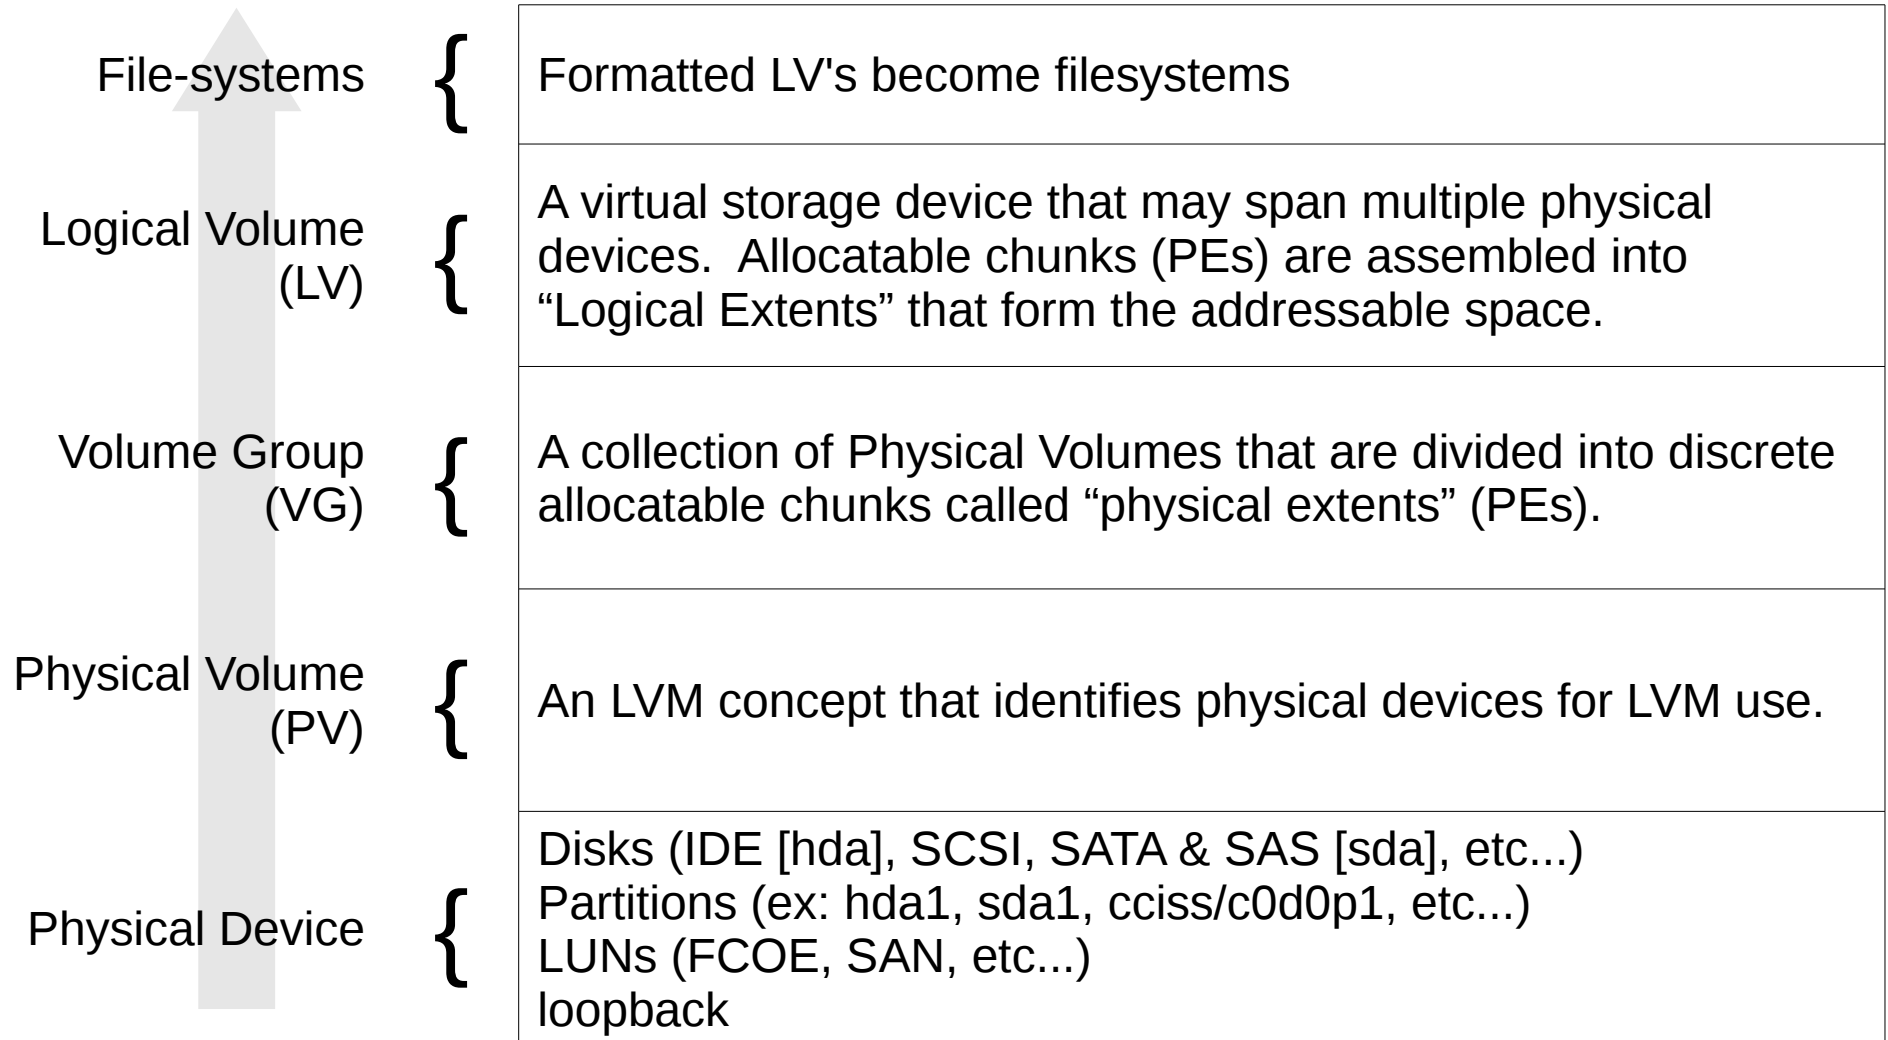
LAB Synopsis

The purpose of this lab is to teach the fundamental concepts of LVM and demonstrate the latest LVM features.

Why use Logical Volume Management?

- Agility
 - Dynamically adjust storage configuration
 - Grow, shrink or relocate your data/filesystems
 - Aggregate or subdivide devices as needed
- Performance
 - Striping across multiple devices
- Fault Tolerance (redundancy & resiliency)
 - RAID 0, 1, 5, 6, 10
 - Snapshots

Building Blocks of Storage Management



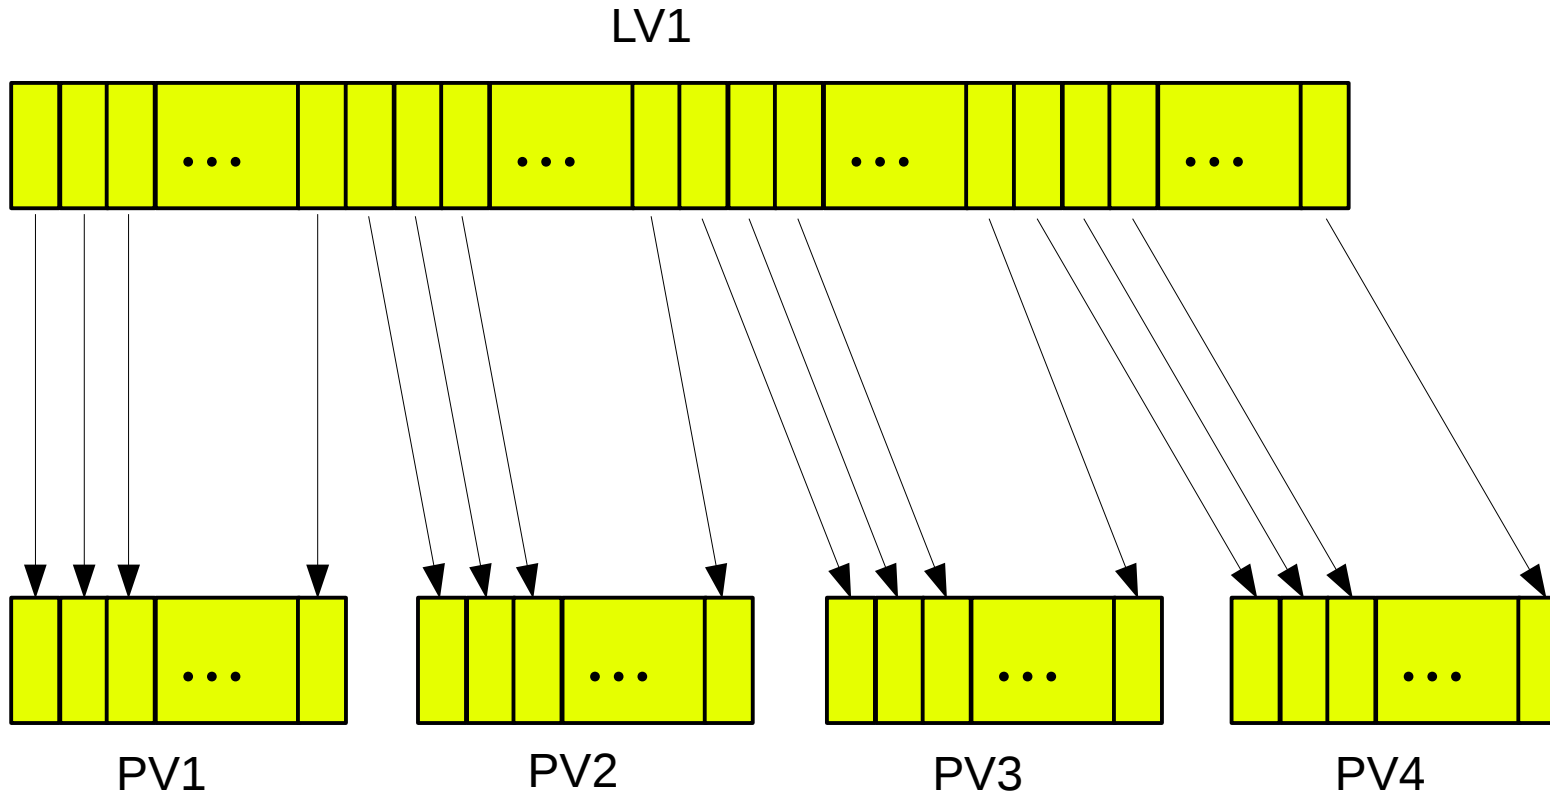
The LVM CLI Toolbox

	Physical Volumes	Volume Groups	Logical Volumes
Core Utilities	pvcreate pvdisplay pvremove pvs pvscan pvmove	vgcreate vgdisplay vgextend vgreduce vgremove vgrename vgs vgscan vgcfgbackup vgcfgrestore	lvconvert lvcreate lvdisplay lvextend lvreduce lvremove lvrename lvresize lvs lvscan
Other Stuff	fdisk parted partprobe multipath smartd		mkfs mount

LVM Linear Volumes

- Summary:
 - A Linear Volume is a virtual device of any size, possibly composed of multiple concatenated physical devices
- Usage:
 - `lvcreate -n <lv_name> -L <size> <VG>`
- Additional Comments:
 - Many options have short & long identifiers
 - '--size' and '-L' are equivalent
 - Options to specify size by # of extents, % free space

Mapping Linear Logical Extents (LE) to Physical Extents (PE)



Lab 1 : Part 1 – Create a Filesystem

Device/Partition

```
fdisk /dev/vdb
```

create partition vdb1 & set type "8E"

```
partprobe
```

Physical Volume

```
pvcreate /dev/vdb1
```

Volume Group

```
vgcreate vg_summit /dev/vdb1
```

Logical Volume

```
lvcreate -n lab1 -L 10M vg_summit
```

Filesystem

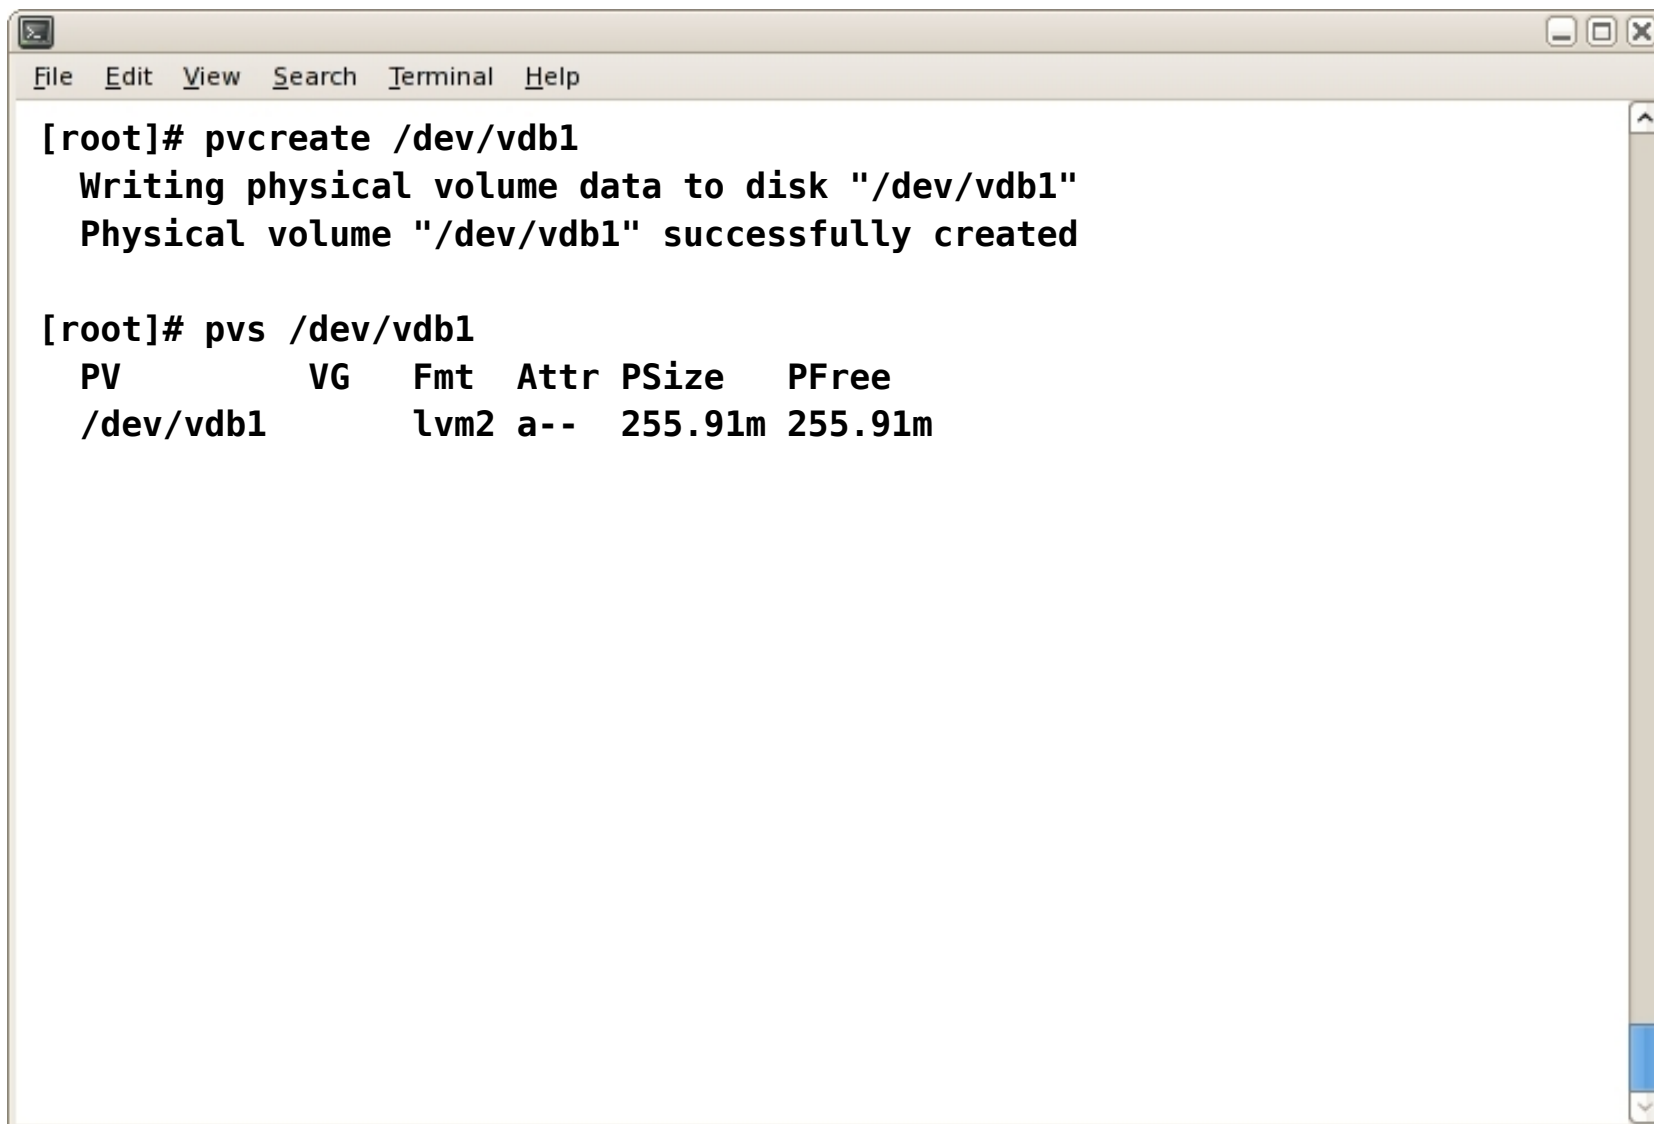
```
mkfs -t ext4 /dev/vg_summit/lab1
```

Mount

```
mkdir -p /mnt/lab1
```

```
mount /dev/vg_summit/lab1 /mnt/lab1
```

Lab 1 : Part 1 – Analyzing pvcreate

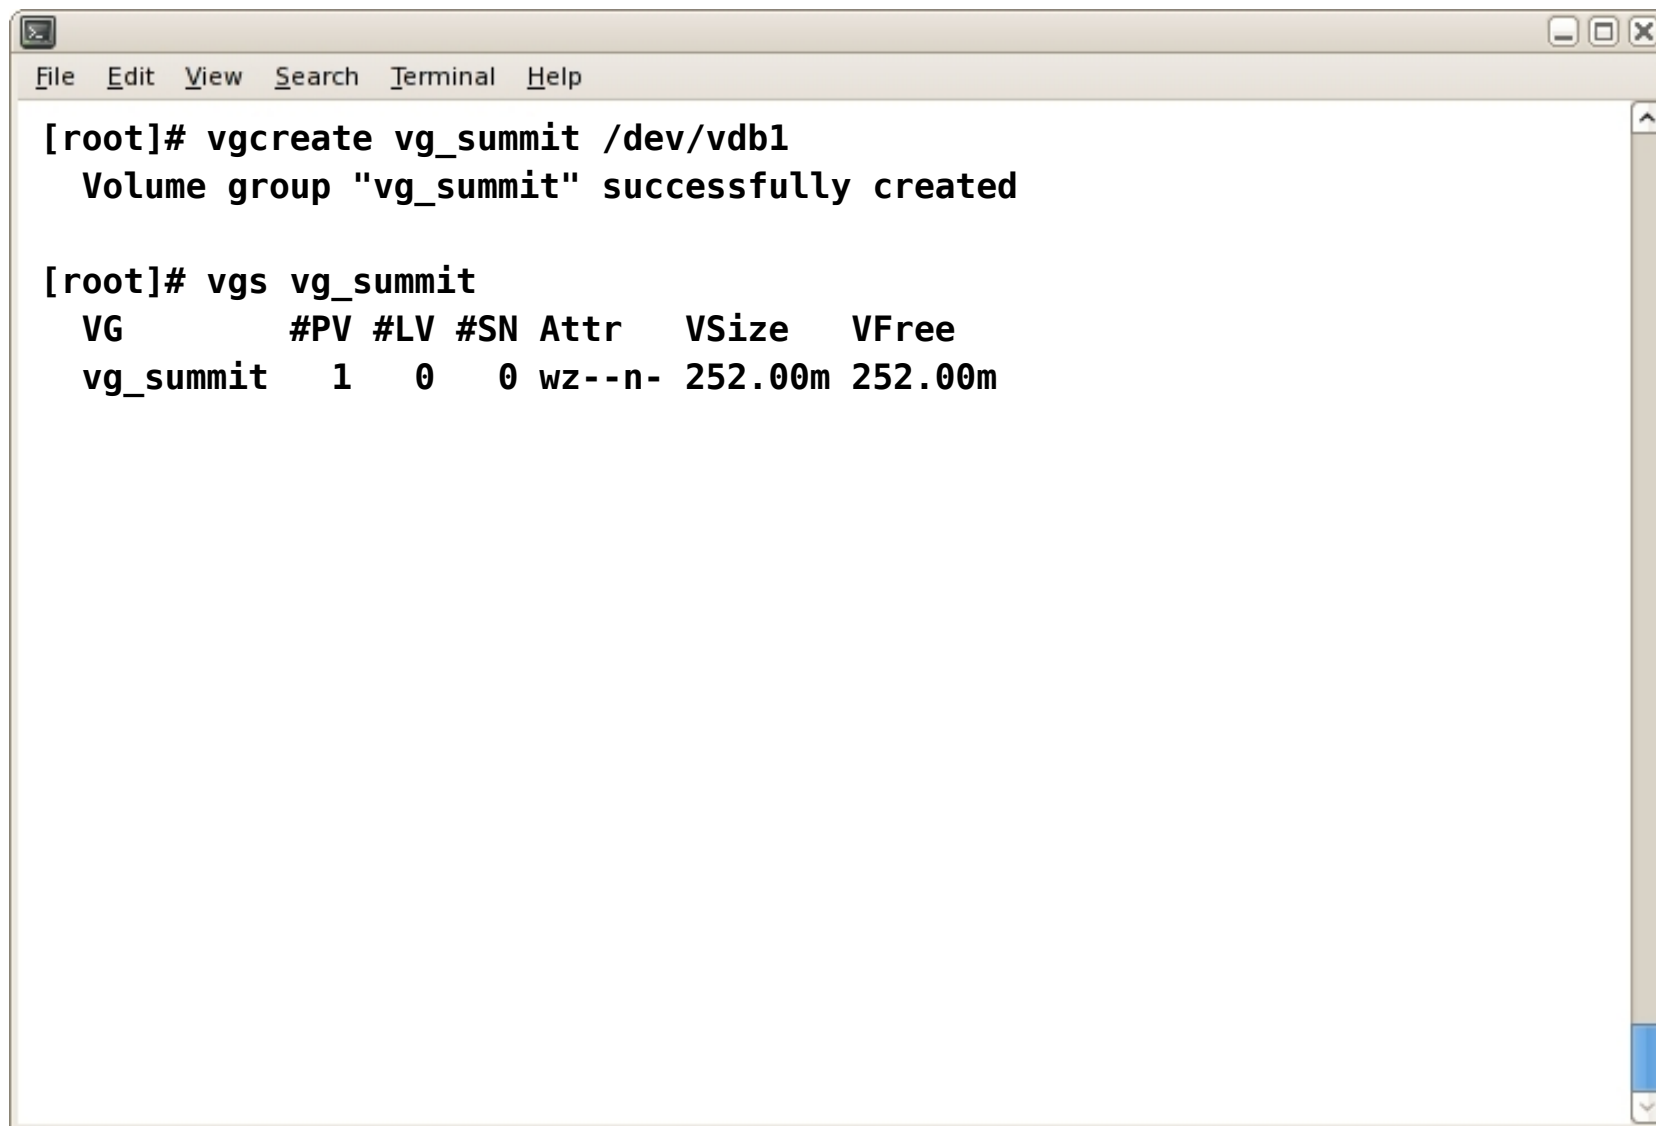


```
File Edit View Search Terminal Help

[root]# pvcreate /dev/vdb1
Writing physical volume data to disk "/dev/vdb1"
Physical volume "/dev/vdb1" successfully created

[root]# pvs /dev/vdb1
PV          VG      Fmt  Attr PSize  PFree
/dev/vdb1   lvm2  a--  255.91m 255.91m
```

Lab 1 : Part 1 – Analyzing vgcreate

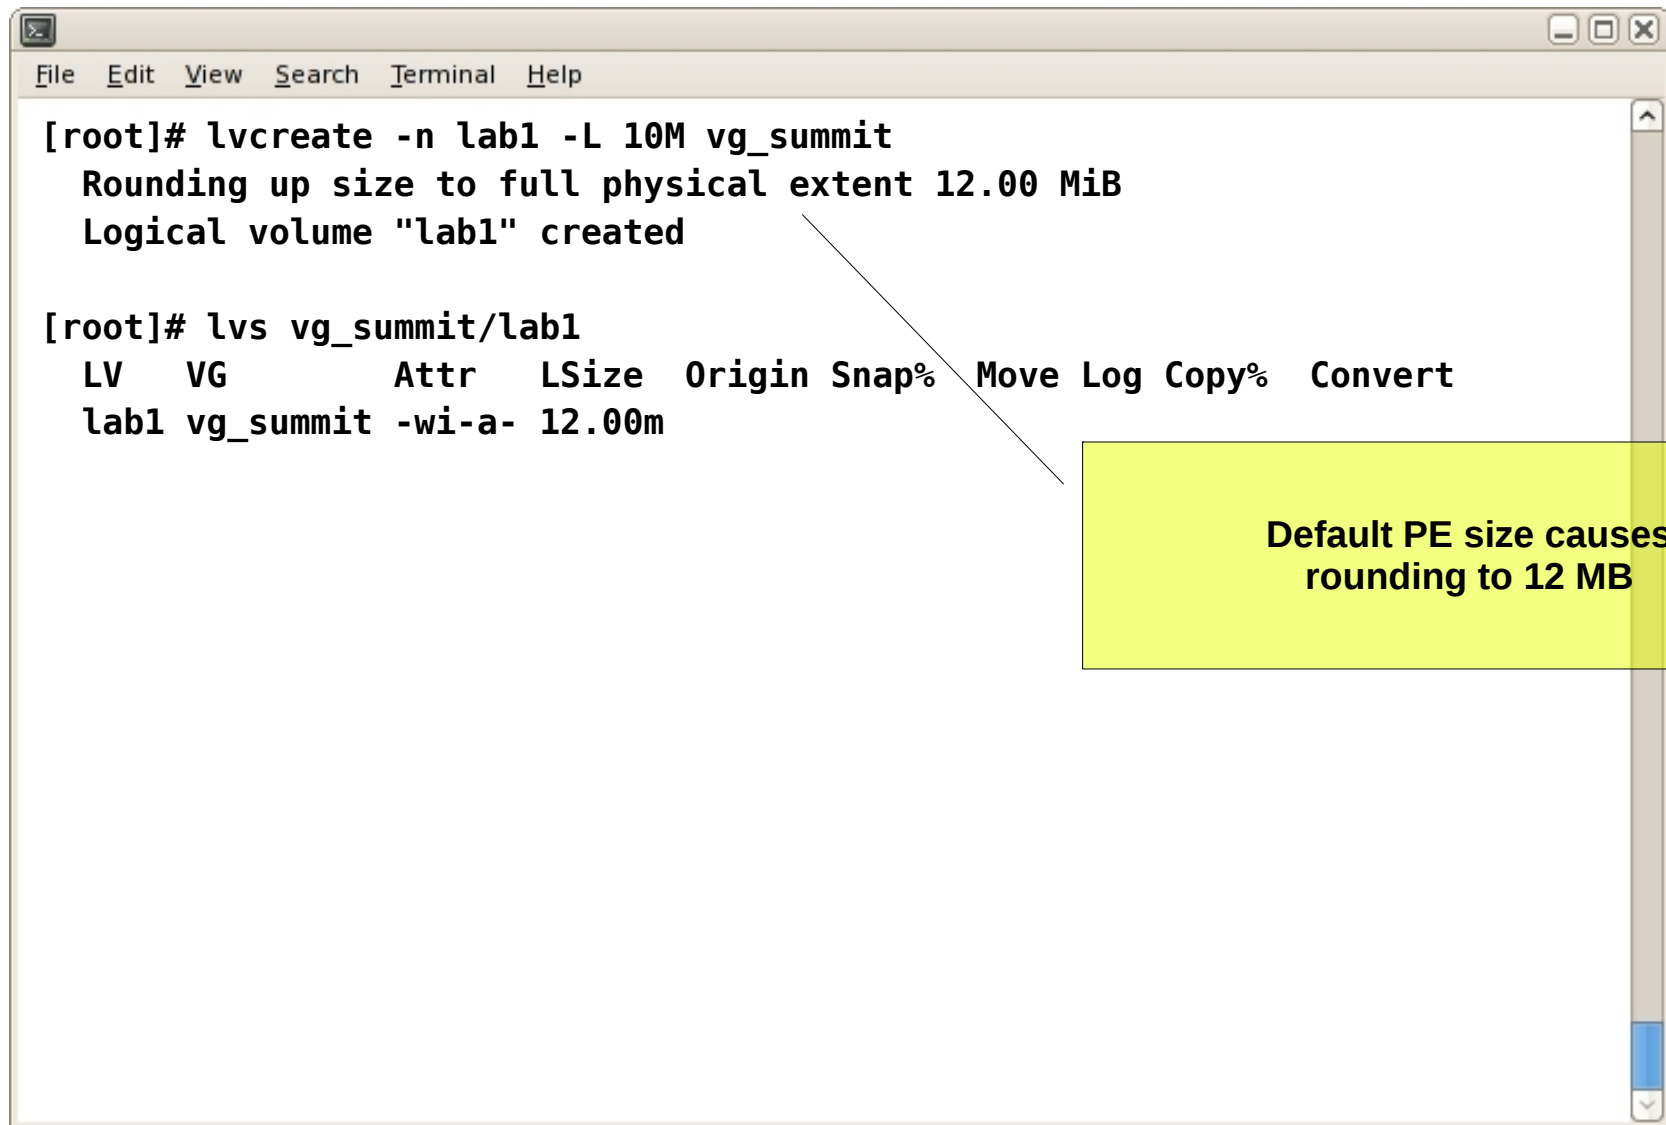


```
File Edit View Search Terminal Help

[root]# vgcreate vg_summit /dev/vdb1
Volume group "vg_summit" successfully created

[root]# vgs vg_summit
VG          #PV #LV #SN Attr   VSize  VFree
vg_summit   1   0   0 wz--n- 252.00m 252.00m
```

Lab 1 : Part 1 – Analyzing lvcreate



```
[root]# lvcreate -n lab1 -L 10M vg_summit
Rounding up size to full physical extent 12.00 MiB
Logical volume "lab1" created

[root]# lvs vg_summit/lab1
LV   VG          Attr   LSize   Origin Snap%   Move Log Copy%  Convert
lab1 vg_summit  -wi-a- 12.00m
```

Default PE size causes rounding to 12 MB

Lab 1 : Part 1 – Analyzing mkfs

```
[root]# mkfs -t ext4 /dev/vg_summit/lab1
mke2fs 1.41.12 (17-May-2010)
Filesystem label=
OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
Stride=0 blocks, Stripe width=0 blocks
3072 inodes, 12288 blocks
614 blocks (5.00%) reserved for the super user
First data block=1
Maximum filesystem blocks=12582912
2 block groups
8192 blocks per group, 8192 fragments per group
1536 inodes per group
Superblock backups stored on blocks:
8193

Writing inode tables: done
Creating journal (1024 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 37 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
```

optional filesystem label

5% could be a lot

fsck every 37 mounts or
180 days

Lab 1 : Part 1 – Analyzing mount

```
File Edit View Search Terminal Help

[root]# mkdir -p /mnt/lab1

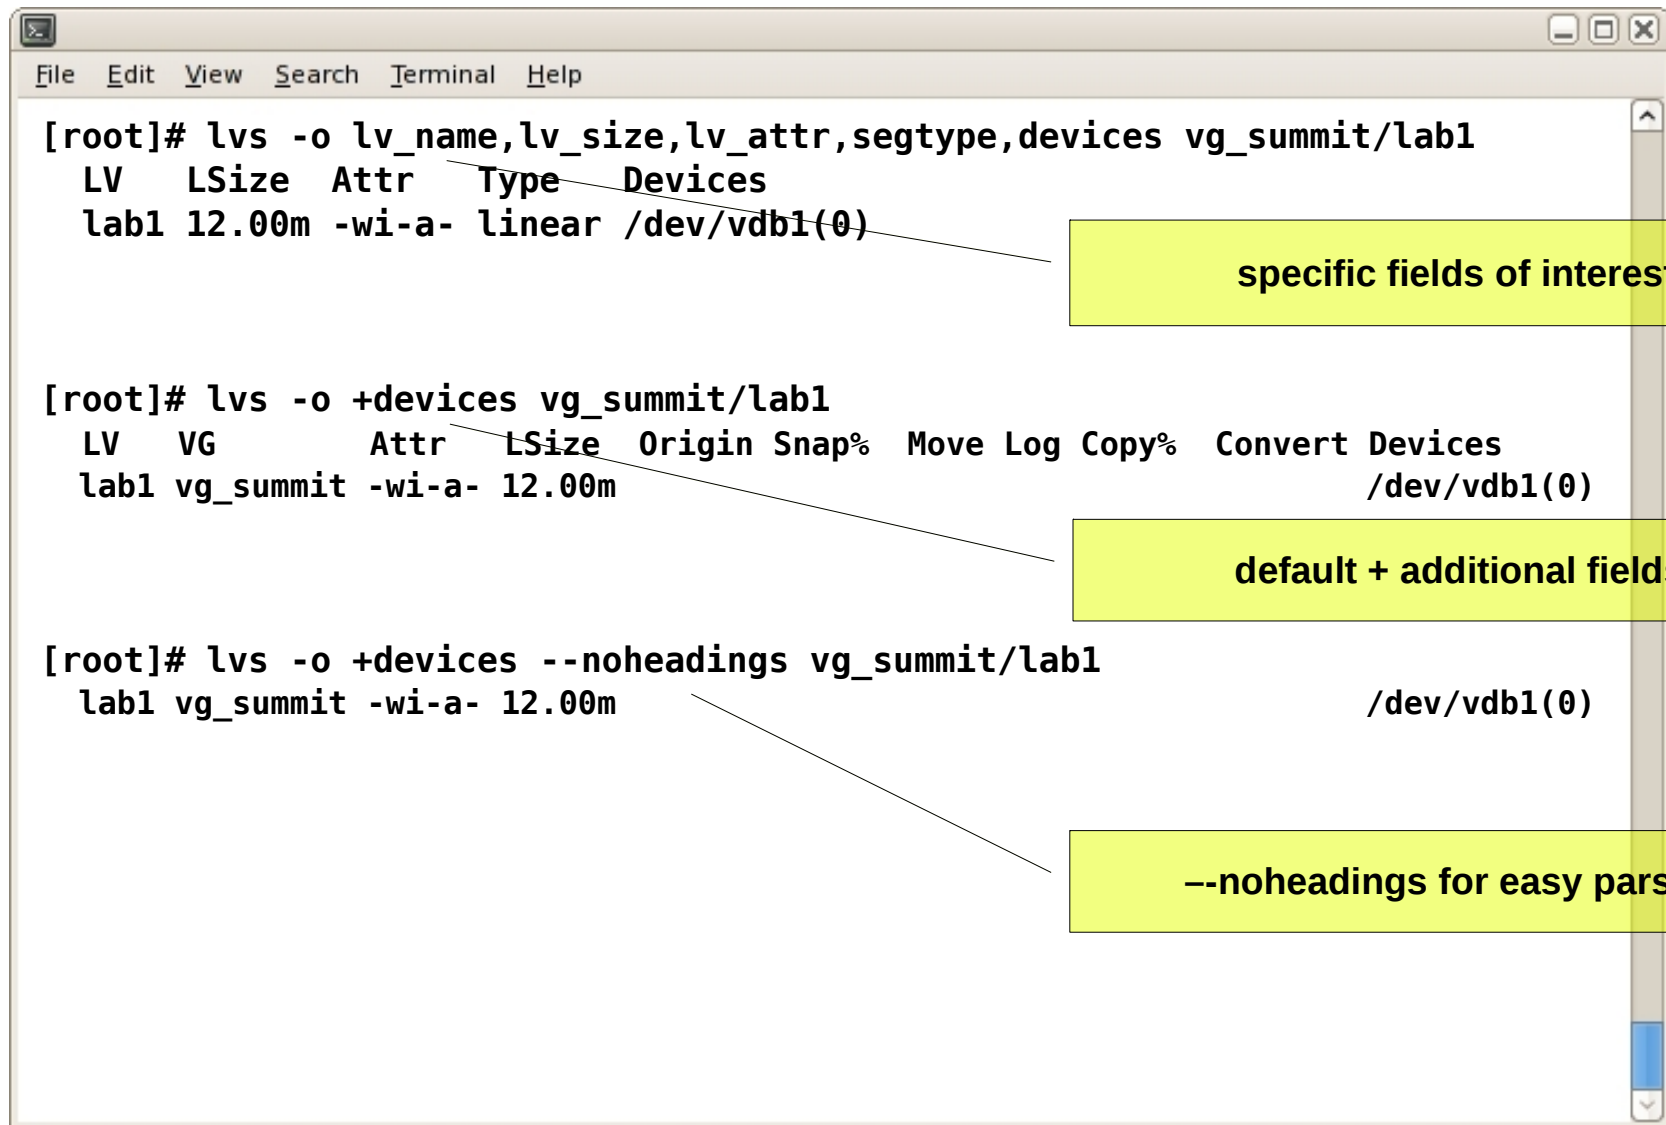
[root]# mount /dev/vg_summit/lab1 /mnt/lab1

[root]# df -h
Filesystem                Size      Used Avail Use% Mounted on
/dev/mapper/vg_desktop-lv_root
                          9.2G     2.1G   6.7G  24% /
tmpfs                      499M     272K   499M   1% /dev/shm
/dev/vda1                  485M      33M   428M   7% /boot
/dev/mapper/vg_summit-lab1
                          12M      1.2M    10M  11% /mnt/lab1
```

use -p to create entire path

/dev/mapper/vg_summit-lab1
vs
/dev/vg_summit/lab1

Lab 1 : Part 1 – additional options & output



```
File Edit View Search Terminal Help

[root]# lvs -o lv_name,lv_size,lv_attr,segtype,devices vg_summit/lab1
LV  LSize Attr  Type  Devices
lab1 12.00m -wi-a- linear /dev/vdb1(0)

[root]# lvs -o +devices vg_summit/lab1
LV  VG      Attr  LSize  Origin Snap%  Move Log Copy%  Convert Devices
lab1 vg_summit -wi-a- 12.00m                               /dev/vdb1(0)

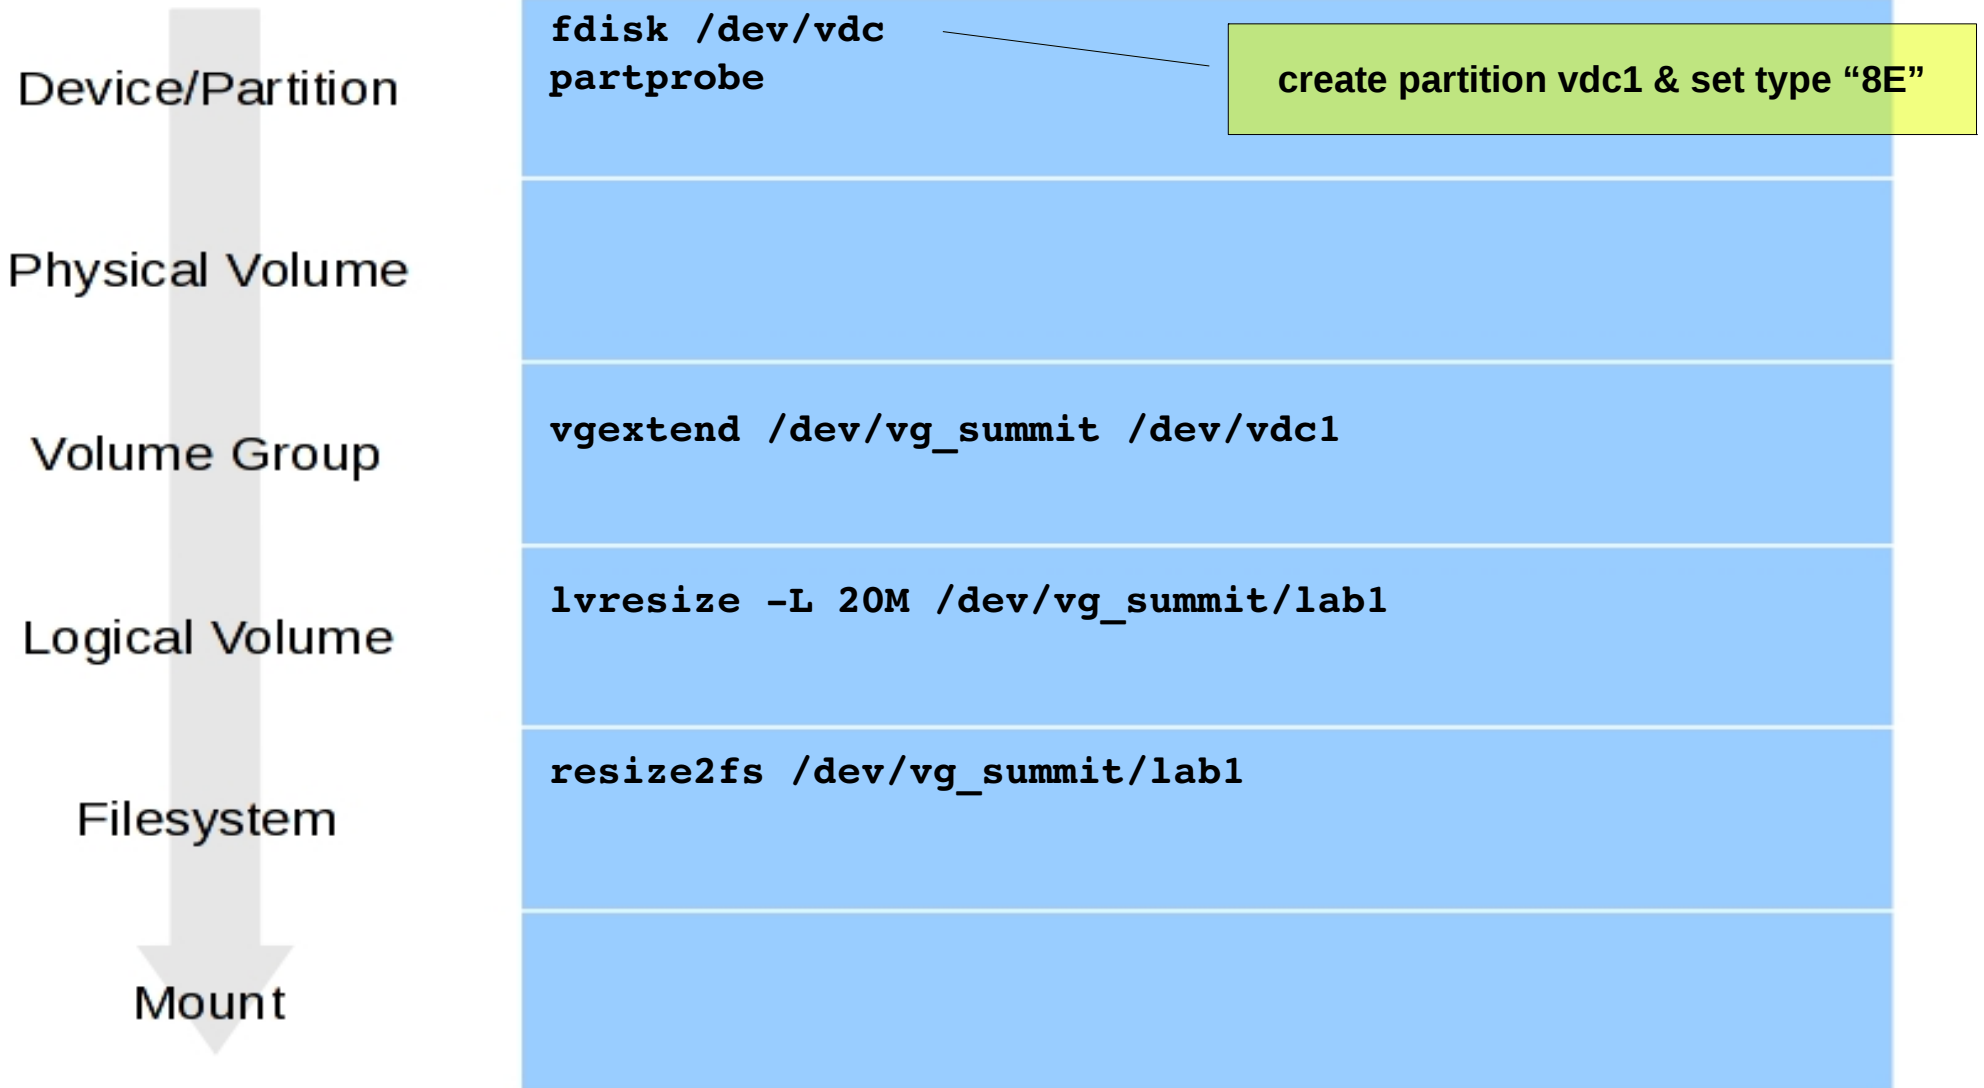
[root]# lvs -o +devices --noheadings vg_summit/lab1
lab1 vg_summit -wi-a- 12.00m                               /dev/vdb1(0)
```

specific fields of interest

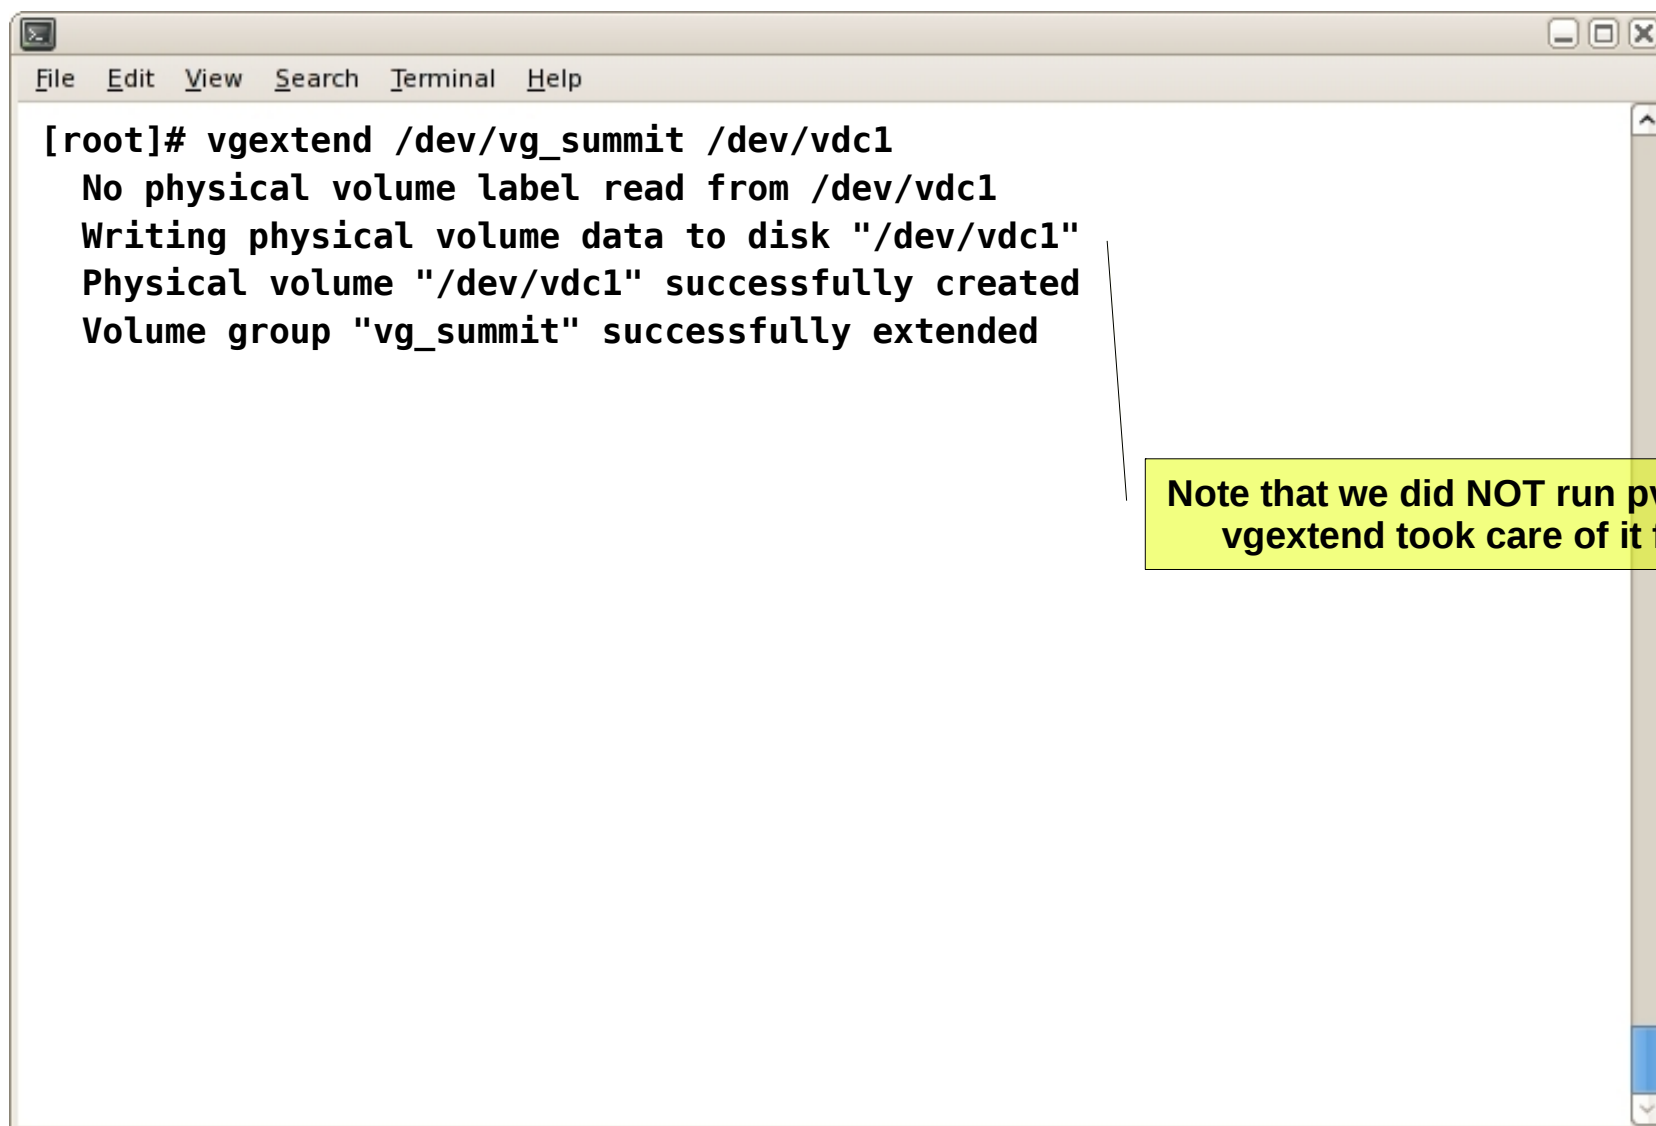
default + additional fields

--noheadings for easy parsing

Lab 1 : Part 2 – Resize a Filesystem



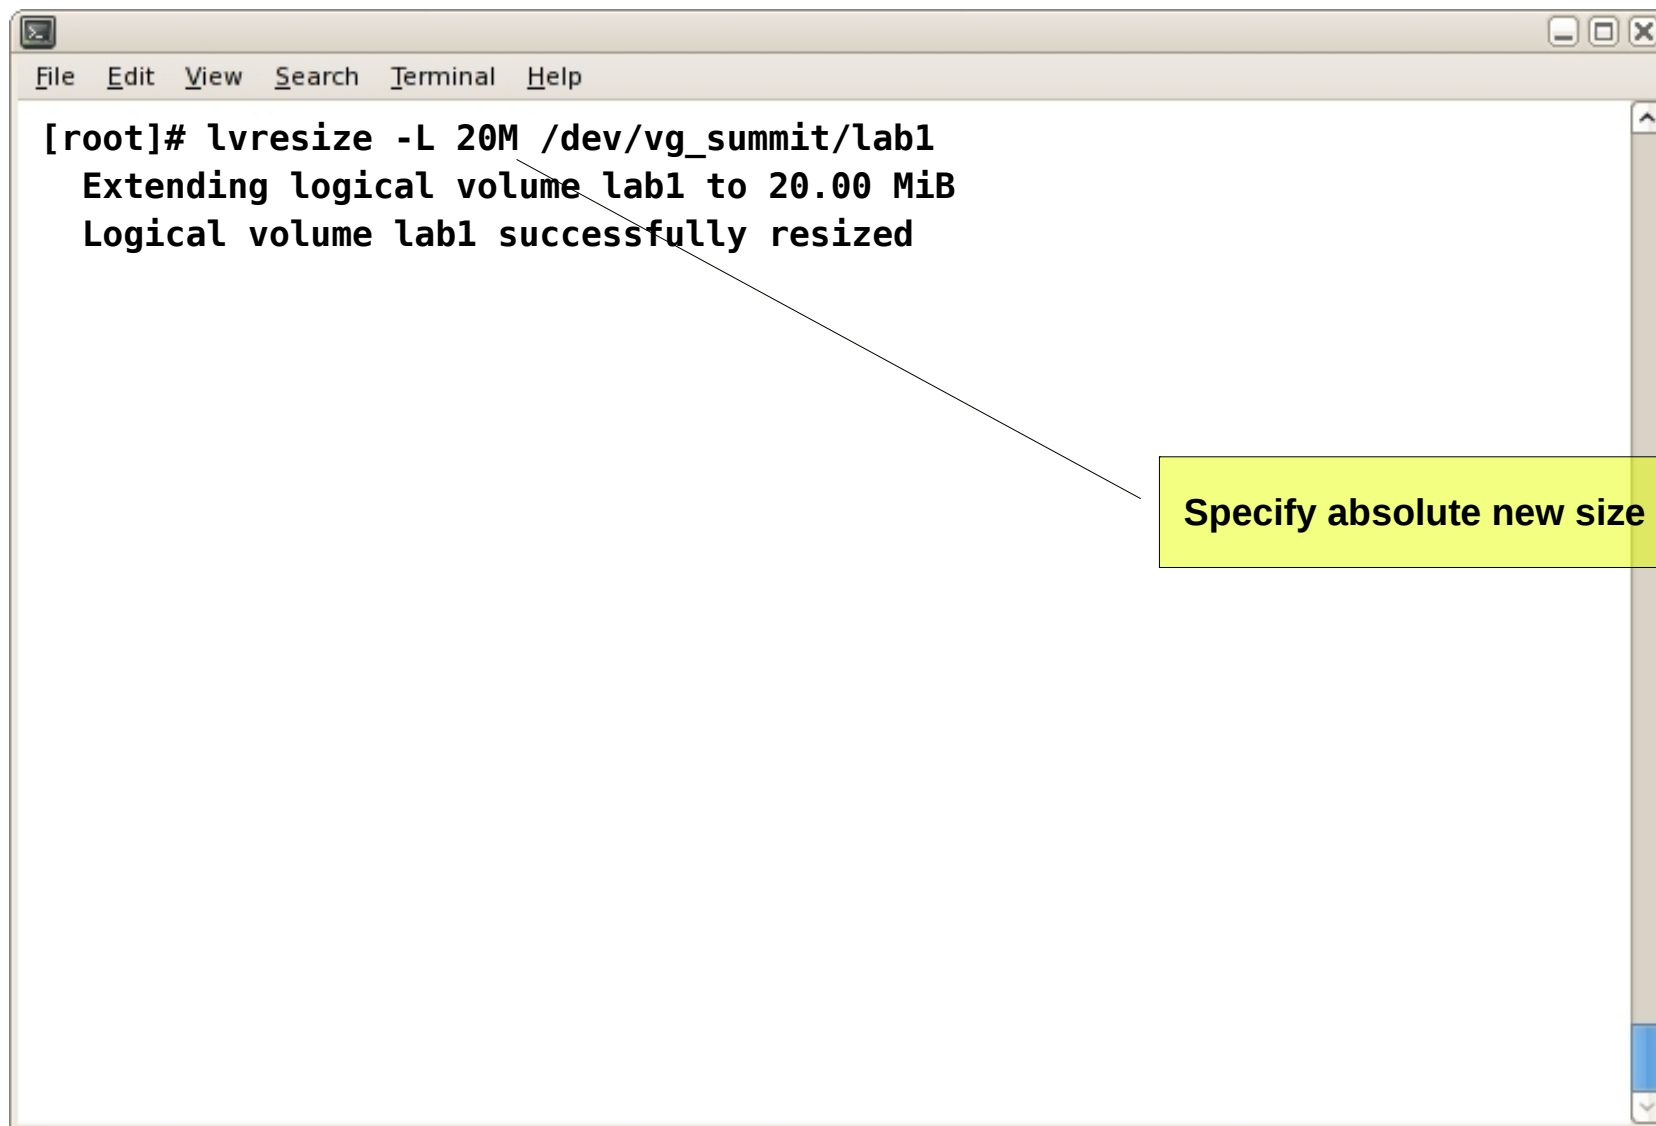
Lab 1 : Part 2 – Analyzing vgextend



```
[root]# vgextend /dev/vg_summit /dev/vdc1
No physical volume label read from /dev/vdc1
Writing physical volume data to disk "/dev/vdc1"
Physical volume "/dev/vdc1" successfully created
Volume group "vg_summit" successfully extended
```

**Note that we did NOT run pvcreate.
vgextend took care of it for us**

Lab 1 : Part 2 – Analyzing lvresize



A terminal window with a menu bar (File, Edit, View, Search, Terminal, Help) and a scroll bar on the right. The terminal output shows the execution of the `lvresize` command. A yellow callout box with a black border points to the `-L 20M` option in the command, containing the text "Specify absolute new size as 20MB".

```
[root]# lvresize -L 20M /dev/vg_summit/lab1
Extending logical volume lab1 to 20.00 MiB
Logical volume lab1 successfully resized
```

Specify absolute new size as 20MB

Lab 1 : Part 2 – Analyzing resize2fs

```
[root]# resize2fs /dev/vg_summit/lab1
resize2fs 1.41.12 (17-May-2010)
Filesystem at /dev/vg_summit/lab1 is mounted on /mnt/lab1; on-line resizing
required
old desc_blocks = 1, new_desc_blocks = 1
Performing an on-line resize of /dev/vg_summit/lab1 to 20480 (1k) blocks.
The filesystem on /dev/vg_summit/lab1 is now 20480 blocks long.

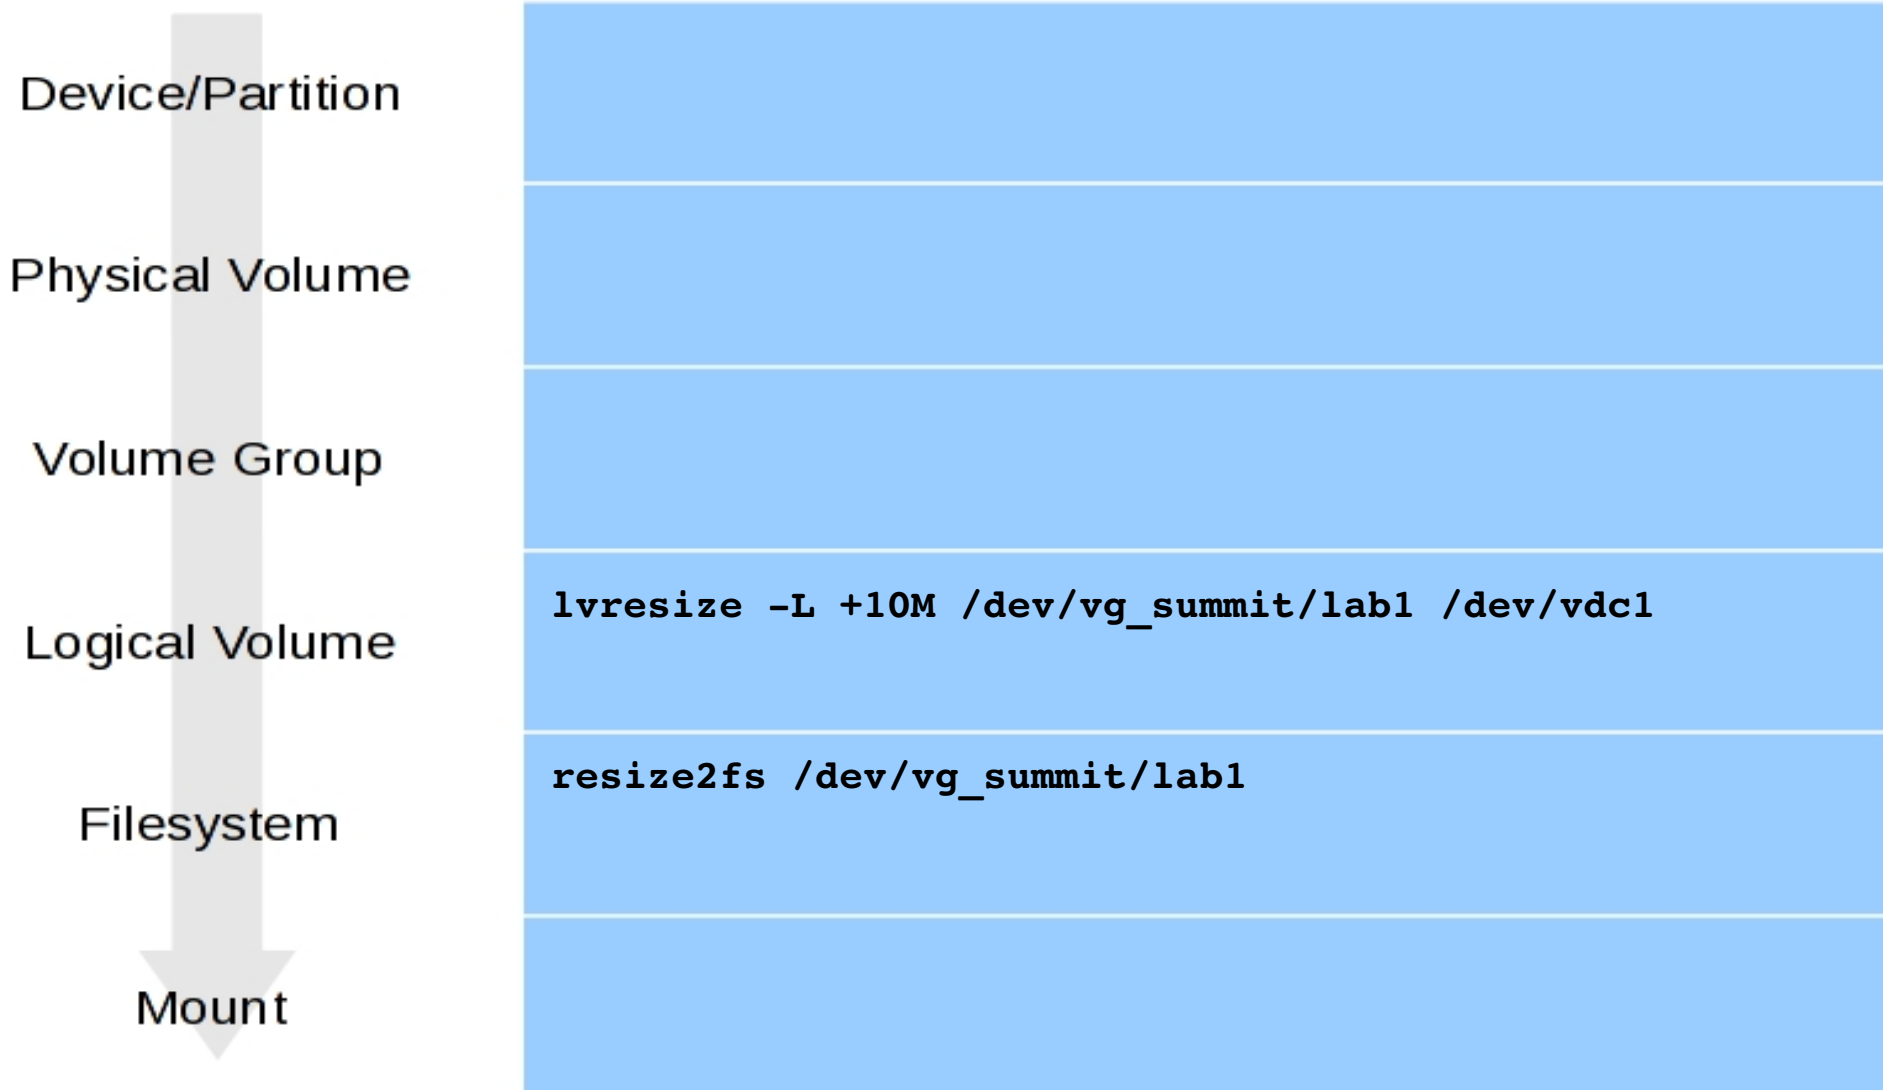
[root]# df -h /mnt/lab1
Filesystem                Size  Used Avail Use% Mounted on
/dev/mapper/vg_summit-lab1 20M   1.2M   18M   6% /mnt/lab1

[root]# lvs -o lv_name,lv_size,lv_attr,segtype,devices vg_summit/lab1
LV  LSize  Attr  Type  Devices
lab1 20.00m -wi-ao linear /dev/vdb1(0)
```

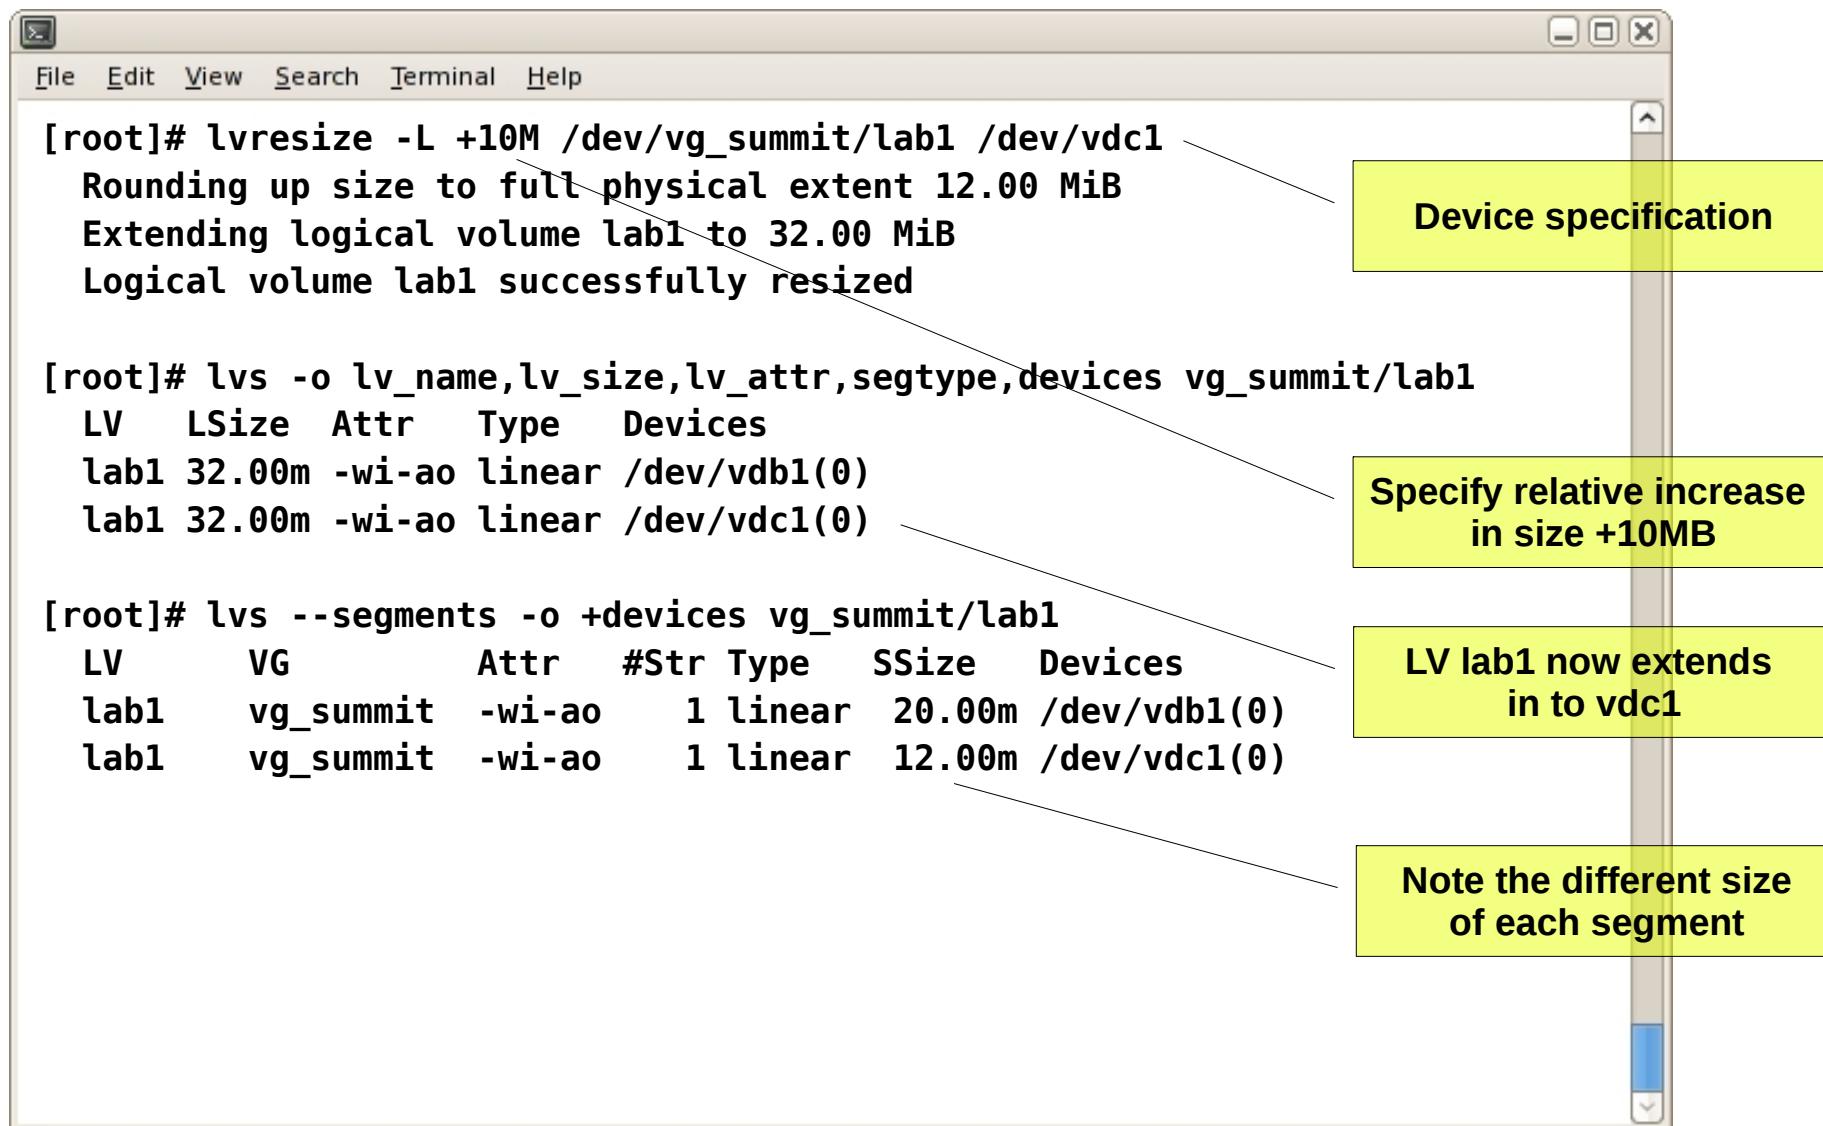
Size is now 20MB

vdb1 still had space, so
resize did not use vdc1

Lab 1 : Part 3 – Resize a Filesystem again



Lab 1 : Part 3 – Analyzing lvresize a second time



```
[root]# lvresize -L +10M /dev/vg_summit/lab1 /dev/vdc1
Rounding up size to full physical extent 12.00 MiB
Extending logical volume lab1 to 32.00 MiB
Logical volume lab1 successfully resized

[root]# lvs -o lv_name,lv_size,lv_attr,segtype,devices vg_summit/lab1
LV  LSize Attr  Type  Devices
lab1 32.00m -wi-ao linear /dev/vdb1(0)
lab1 32.00m -wi-ao linear /dev/vdc1(0)

[root]# lvs --segments -o +devices vg_summit/lab1
LV      VG          Attr  #Str Type  SSize  Devices
lab1    vg_summit  -wi-ao  1 linear 20.00m /dev/vdb1(0)
lab1    vg_summit  -wi-ao  1 linear 12.00m /dev/vdc1(0)
```

Device specification

Specify relative increase in size +10MB

LV lab1 now extends in to vdc1

Note the different size of each segment

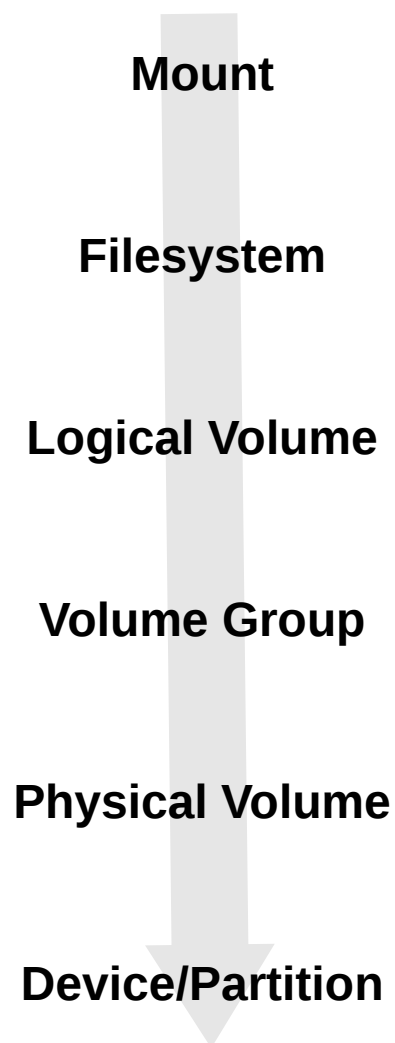
Lab 1 : Part 3 – Analyzing resize2fs a second time

```
[root]# resize2fs /dev/vg_summit/lab1
resize2fs 1.41.12 (17-May-2010)
Filesystem at /dev/vg_summit/lab1 is mounted on /mnt/lab1; on-line resizing
required
old desc_blocks = 1, new_desc_blocks = 1
Performing an on-line resize of /dev/vg_summit/lab1 to 32768 (1k) blocks.
The filesystem on /dev/vg_summit/lab1 is now 32768 blocks long.

[root]# df -h /mnt/lab1
Filesystem                Size      Used Avail Use% Mounted on
/dev/mapper/vg_summit-lab1  32M    1.2M   29M   4% /mnt/lab1
```

Size is now 32MB

Lab 1 : Part 4 – Removal & Cleanup



```
umount /mnt/lab1
```

Optional : zero out data

```
dd if=/dev/zero of=/dev/vg_summit/lab1
```

```
lvremove -f vg_summit/lab1
```

```
vgremove -f vg_summit
```

```
pvremove -f /dev/vd{b,c}1
```

RED HAT
SUMMIT

Logical Volume Management

Striping and mirroring

LVM Stripes

- Summary:
 - Striping (RAID0) aggregates storage devices, increases performance BUT provides no redundancy.
- Usage:
 - `lvcreate -i <stripe_count> -n <lv_name> -L <size> <vg>`
- Additional Comments:
 - '-i' and '--stripes' are equivalent
 - Set the stripe size: '-l/--stripesize <StripeSize>'

Lab 2 : Part 1 – Create a Striped LV

Device/Partition

Physical Volume

Volume Group

Logical Volume

Filesystem

Mount

```
vgcreate vg_summit /dev/vd{b,c,d}1
```

```
lvcreate -i 3 -n lab2 -L 100M vg_summit
```

```
mkfs -t ext4 /dev/vg_summit/lab2
```

```
mkdir -p /mnt/lab2
```

```
mount /dev/vg_summit/lab2 /mnt/lab2
```

Lab 2 : Part 1 – Analyzing vgcreate

```
[root]# vgcreate vg_summit /dev/vd{b,c,d}1
No physical volume label read from /dev/vdb1
No physical volume label read from /dev/vdc1
No physical volume label read from /dev/vdd1
Writing physical volume data to disk "/dev/vdb1"
Physical volume "/dev/vdb1" successfully created
Writing physical volume data to disk "/dev/vdc1"
Physical volume "/dev/vdc1" successfully created
Writing physical volume data to disk "/dev/vdd1"
Physical volume "/dev/vdd1" successfully created
Volume group "vg_summit" successfully created
```

Note that we did NOT run pvcreate.
vgcreate took care of it for us

Lab 2 : Part 1 – Analyzing lvcreate

```
File Edit View Search Terminal Help

[root]# lvcreate -i 3 -n lab2 -L 100M vg_summit
Using default stripesize 64.00 KiB
Rounding size (25 extents) up to stripe boundary size (27 extents)
Logical volume "lab2" created

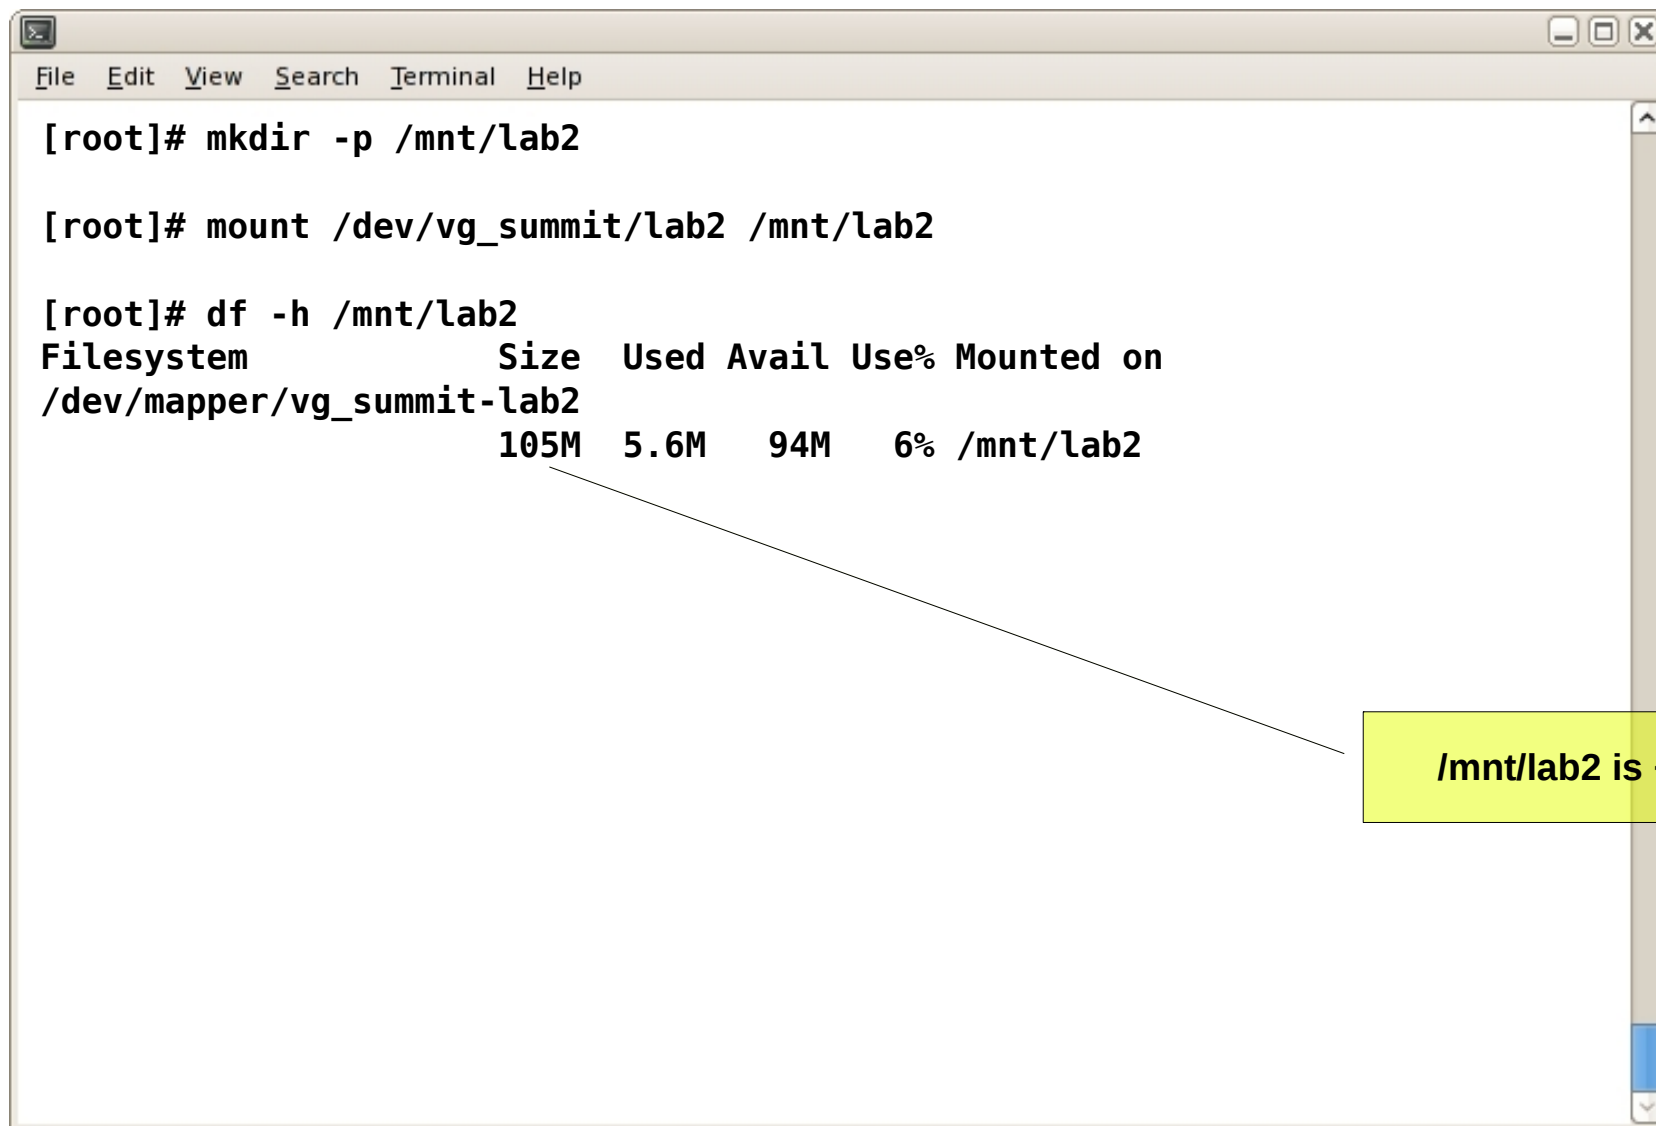
[root]# lvs -o lv_name,lv_size,lv_attr,segtype,devices vg_summit/lab2
LV  LSize  Attr  Type  Devices
lab2 108.00m -wi-a- striped /dev/vdb1(0),/dev/vdc1(0),/dev/vdd1(0)
```

We didn't specify **-I** thus default stripesize

lab2 now stripes across vdb1, vdc1,vdd1

lab2 is striped

Lab 2 : Part 1 – Analyzing mount



```
File Edit View Search Terminal Help

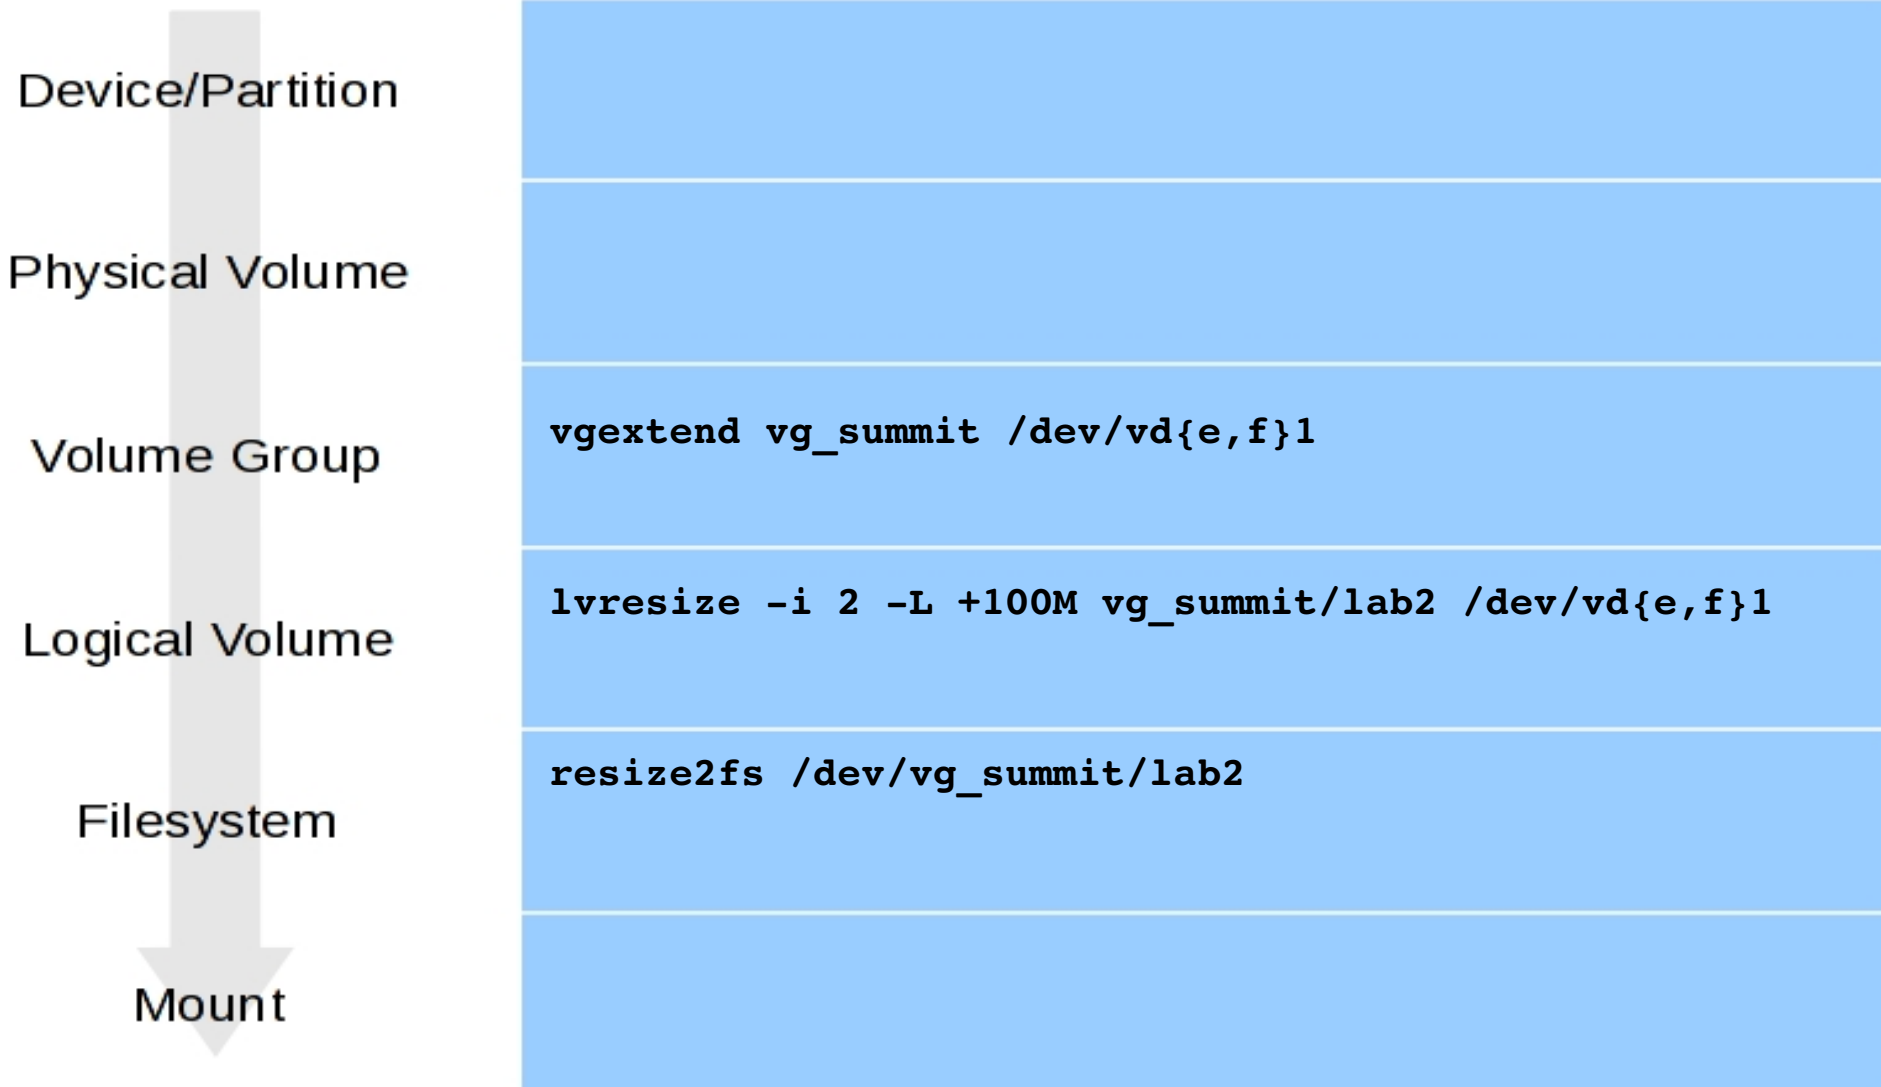
[root]# mkdir -p /mnt/lab2

[root]# mount /dev/vg_summit/lab2 /mnt/lab2

[root]# df -h /mnt/lab2
Filesystem                Size  Used Avail Use% Mounted on
/dev/mapper/vg_summit-lab2
                          105M  5.6M   94M   6% /mnt/lab2
```

/mnt/lab2 is ~100M

Lab 2 : Part 2 – Resize a Striped LV



Lab 2 : Part 2 – Analyzing vgcreate

```
[root]# vgextend vg_summit /dev/vd{e,f}1
No physical volume label read from /dev/vde1
No physical volume label read from /dev/vdf1
Writing physical volume data to disk "/dev/vde1"
Physical volume "/dev/vde1" successfully created
Writing physical volume data to disk "/dev/vdf1"
Physical volume "/dev/vdf1" successfully created
Volume group "vg_summit" successfully extended
```

Note that we did NOT run pvcreate.
vgextend took care of it for us

Lab 2 : Part 2 – Analyzing lvresize

```
[root]# lvresize -i 2 -L +100M vg_summit/lab2 /dev/vd{e,f}1
Using stripesize of last segment 64.00 KiB
Rounding size (52 extents) down to stripe boundary size for segment (53 extents)
Extending logical volume lab2 to 212.00 MiB
Logical volume lab2 successfully resized

[root]# lvs -o lv_name,lv_size,lv_attr,segtype,devices vg_summit/lab2
LV  LSize  Attr  Type  Devices
lab2 212.00m -wi-ao striped /dev/vdb1(0),/dev/vdc1(0),/dev/vdd1(0)
lab2 212.00m -wi-ao striped /dev/vde1(0),/dev/vdf1(0)

[root]# lvs --segments -o +devices vg_summit
LV  VG      Attr  #Str Type  SSize  Devices
lab2 vg_summit -wi-ao 3 striped 108.00m /dev/vdb1(0),/dev/vdc1(0),/dev/vdd1(0)
lab2 vg_summit -wi-ao 2 striped 104.00m /dev/vde1(0),/dev/vdf1(0)
```

lab2 segment 1 has 3 stripes

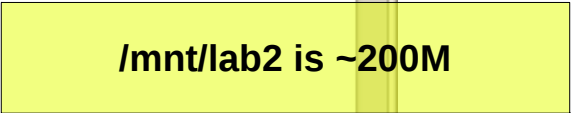
lab2 segment 2 has 2 stripes

Lab 2 : Part 2 – Analyzing resize2fs

```
File Edit View Search Terminal Help

[root]# resize2fs /dev/vg_summit/lab2
resize2fs 1.41.12 (17-May-2010)
Filesystem at /dev/vg_summit/lab2 is mounted on /mnt/lab2; on-line resizing
required
old desc_blocks = 1, new_desc_blocks = 1
Performing an on-line resize of /dev/vg_summit/lab2 to 208896 (1k) blocks.
The filesystem on /dev/vg_summit/lab2 is now 208896 blocks long.

[root]# df -h /mnt/lab2
Filesystem                Size      Used Avail Use% Mounted on
/dev/mapper/vg_summit-lab2 206M    5.8M   190M   3% /mnt/lab2
```



LVM Mirrors

- Summary:
 - LVM mirrors can maintain redundant copies of your data for device fault tolerance.
- Usage:
 - `lvcreate -m <extra_copy_count> -n <lv_name> -L <size> <vg>`
- Additional Comments:
 - Prevent initial synchronization: `'--nosync'`
 - Specify the log type: `'--mirrorlog <core|disk>'`

Lab 3 : Part 1 – Create a Mirrored LV

Device/Partition

Physical Volume

Volume Group

Logical Volume

Filesystem

Mount

```
lvcreate -m 1 -n lab3 -L 20M vg_summit
```

```
mkfs -t ext4 /dev/vg_summit/lab3
```

```
mkdir -p /mnt/lab3
```

```
mount /dev/vg_summit/lab3 /mnt/lab3
```

Lab 3 : Part 1 – Analyzing lvcreate

```
File Edit View Search Terminal Help

[root]# lvcreate -m 1 -n lab3 -L 20M vg_summit
Logical volume "lab3" created

[root]# lvs -o lv_name,lv_size,lv_attr,segtype,devices vg_summit/lab3
LV      LSize  Attr  Type   Devices
lab3    20.00m mwi-ao mirror lab3_mimage_0(0),lab3_mimage_1(0)

[root]# lvs -a -o lv_name,lv_size,lv_attr,segtype,devices vg_summit
LV              LSize  Attr  Type   Devices
lab2            212.00m -wi-ao striped /dev/vdb1(0),/dev/vdc1(0),/dev/vdd1(0)
lab2            212.00m -wi-ao striped /dev/vde1(0),/dev/vdf1(0)
lab3            20.00m mwi-ao mirror  lab3_mimage_0(0),lab3_mimage_1(0)
[lab3_mimage_0] 20.00m iwi-ao linear  /dev/vdb1(9)
[lab3_mimage_1] 20.00m iwi-ao linear  /dev/vdc1(9)
[lab3_mlog]     4.00m lwi-ao linear  /dev/vdf1(12)
```

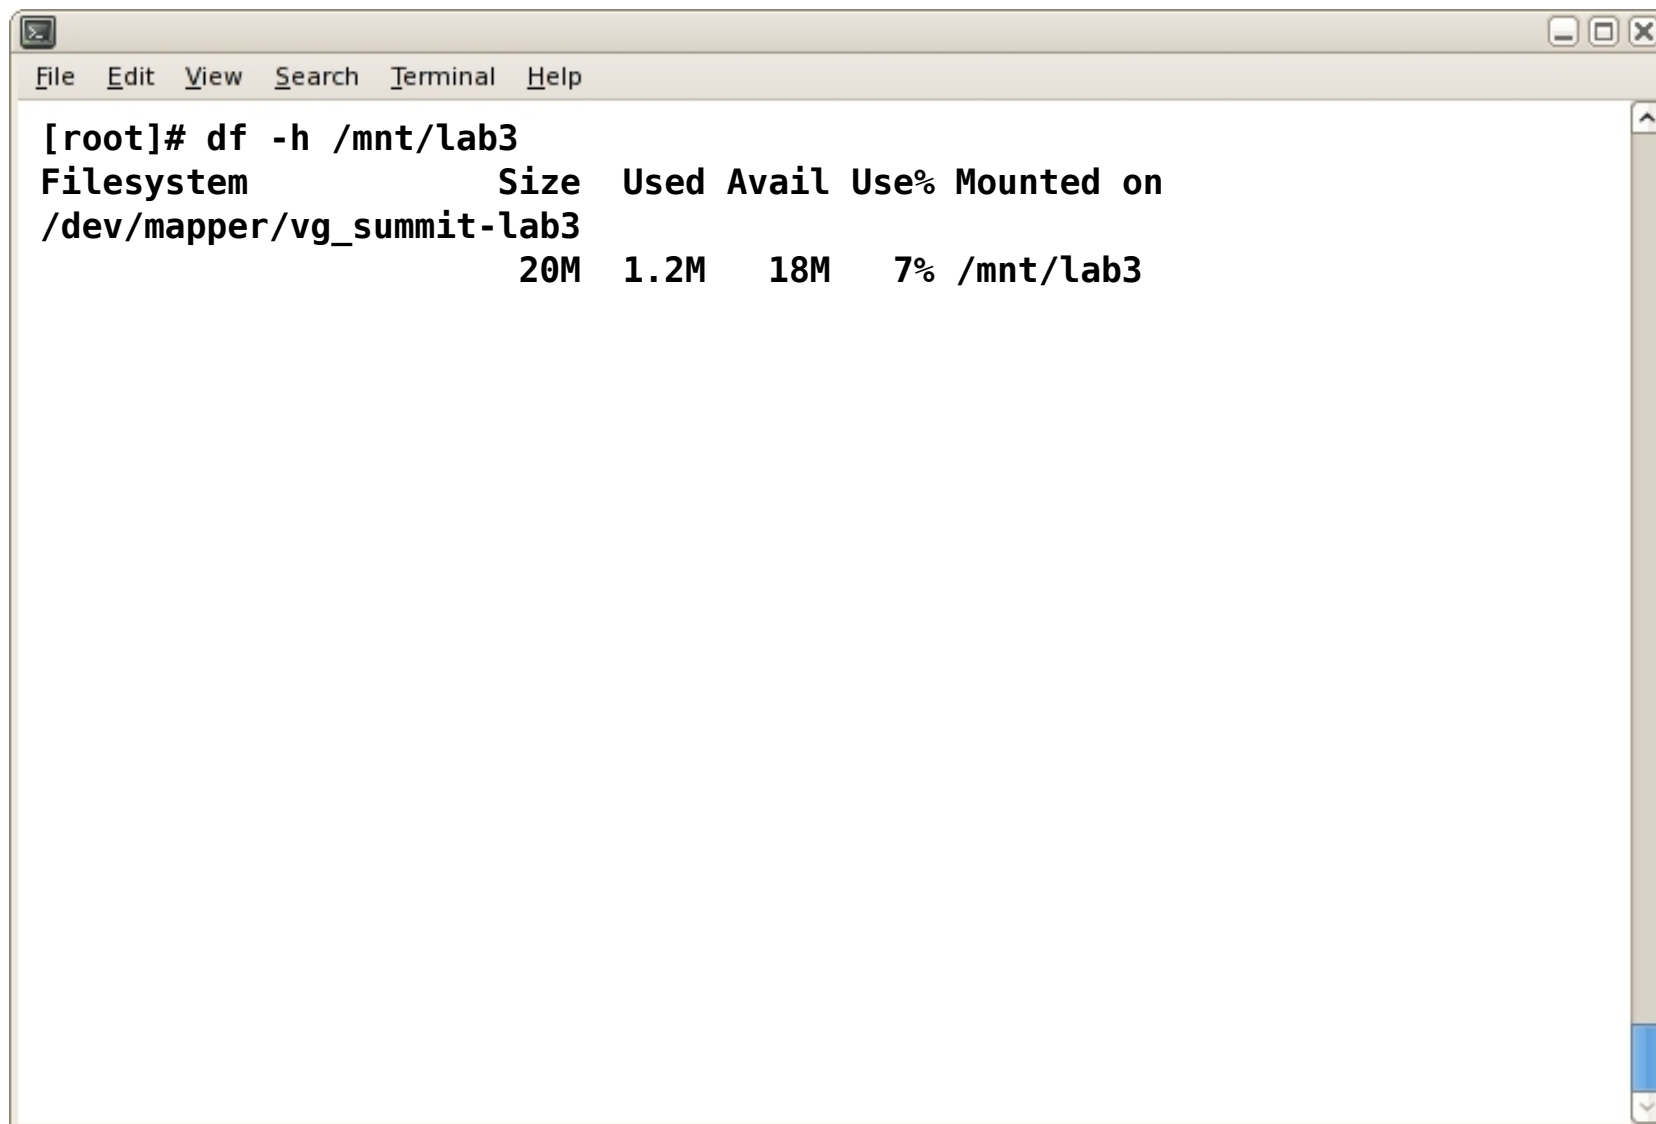
lab3 with 1 additional copy (mirror)

segment is a mirror

'm' attribute indicates mirror
'i' attribute indicates image
'l' attribute indicates log

Subordinate devices in '[']' display as a result of '-a'

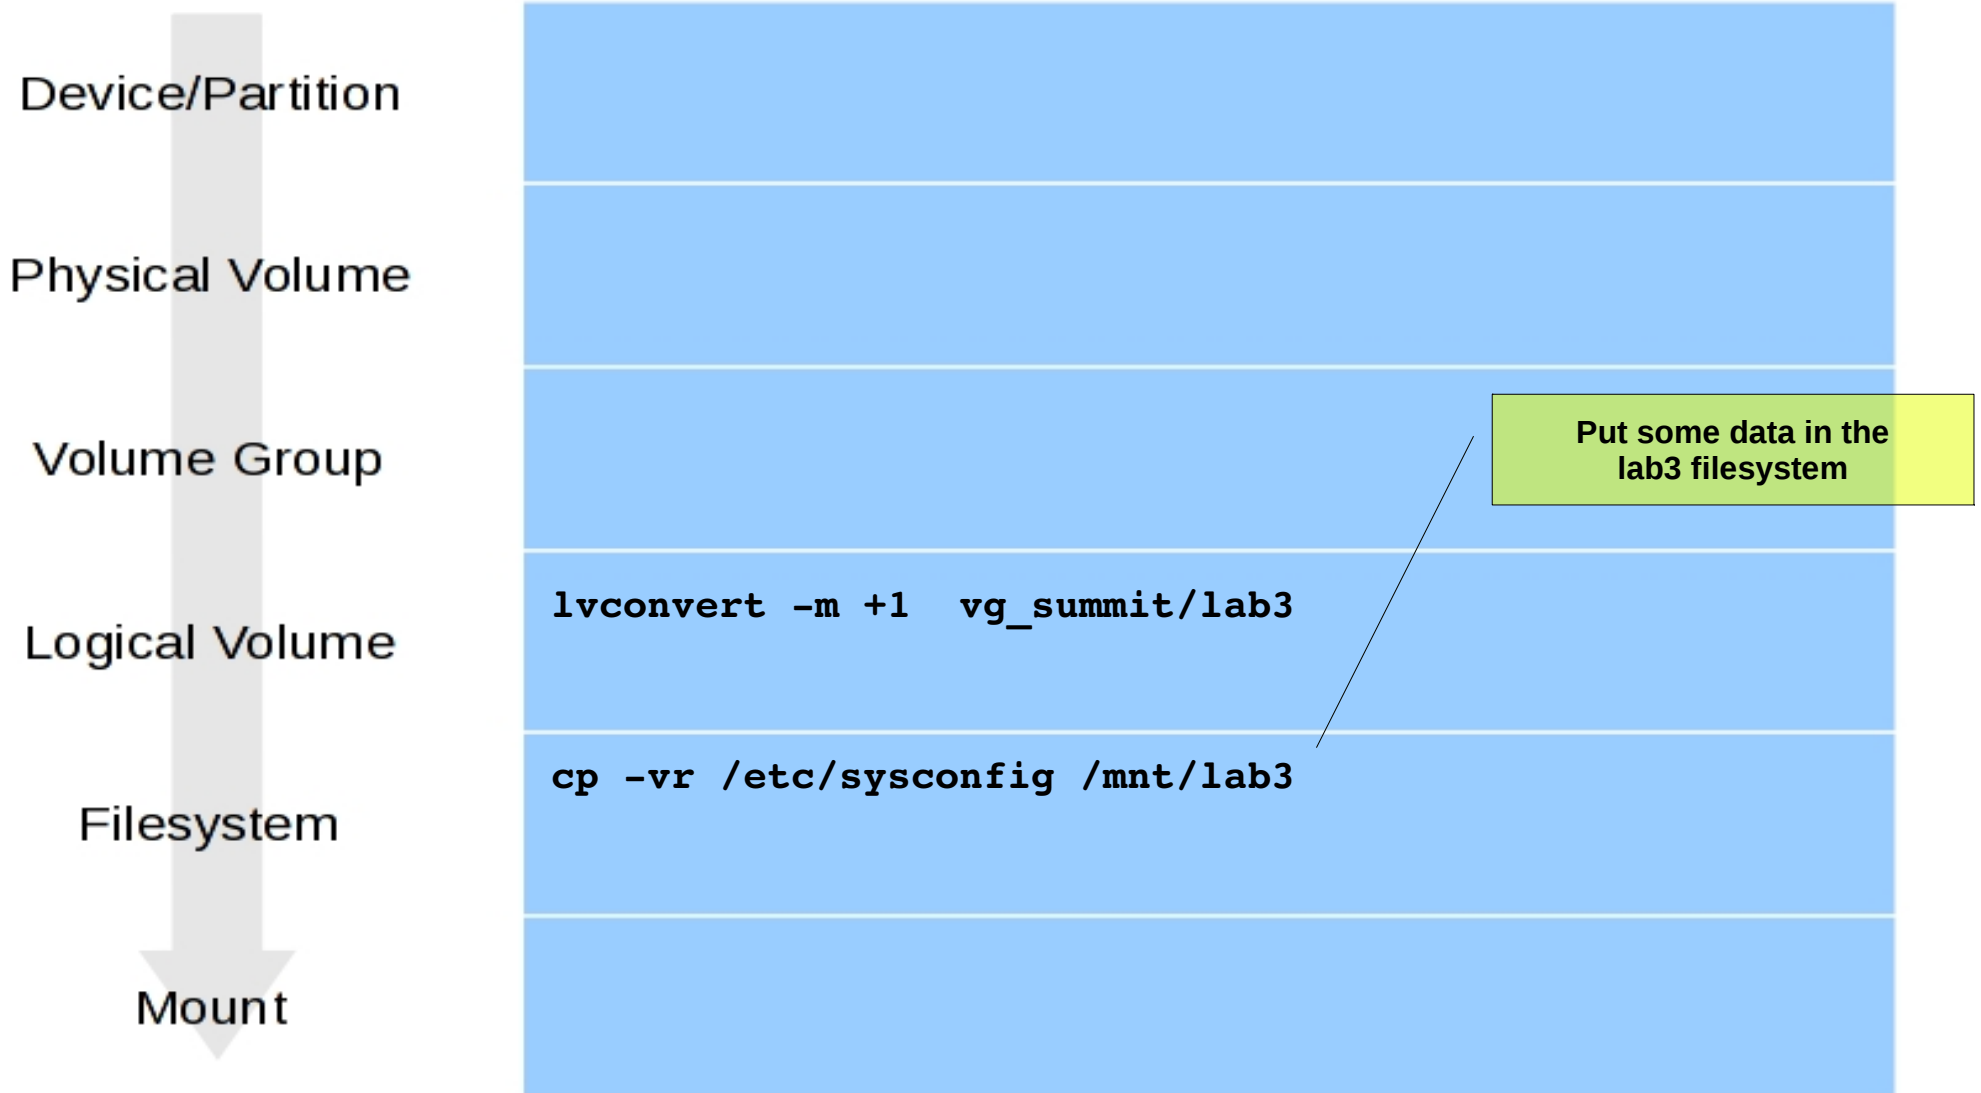
Lab 3 : Part 1 – Analyzing the filesystem



A terminal window with a menu bar (File, Edit, View, Search, Terminal, Help) and a scroll bar on the right. The terminal shows the command `[root]# df -h /mnt/lab3` and its output:

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/mapper/vg_summit-lab3	20M	1.2M	18M	7%	/mnt/lab3

Lab 3 : Part 3 – Upconvert a Mirrored LV (add another leg)



Lab 3 : Part 2 – Analyzing lvconvert

```
[root]# lvconvert -m +1 vg_summit/lab3
vg_summit/lab3: Converted: 40.0%
vg_summit/lab3: Converted: 100.0%
Logical volume lab3 converted.

[root]# lvs -o lv_name,lv_size,lv_attr,segtype,devices vg_summit/lab3
LV   LSize Attr   Type   Devices
lab3 20.00m mwi-ao mirror lab3_mimage_0(0),lab3_mimage_1(0),lab3_mimage_2(0)
```

add 1 additional copy (mirror) to lab3

There are 3 copies of the data

segment is a mirror

'm' attribute also indicates mirror

LVM Procedure - Mirror Split

- Summary:
 - Splitting an image from a mirror gives point-in-time copy of a mirrored LV.
- Usage:
 - `lvconvert --splitmirror 1 -n <new_lv_name> <vg>/<lv>`
- Additional Comments:
 - Split Mirror requires a name argument
 - Mirror must have completed synchronization
 - Be sure to 'sync' before splitting a mirror

Lab 3 : Part 3 – Split a Mirrored LV

Device/Partition

Physical Volume

Volume Group

Logical Volume

Filesystem

Mount

```
sync
```

```
lvconvert --splitmirrors 1 -n lab3_split vg_summit/lab3
```

```
mkdir -p /mnt/lab3_split
```

```
mount /dev/vg_summit/lab3_split /mnt/lab3_split
```

Lab 3 : Part 3 – Analyzing lvconvert

```
[root]# lvconvert --splitmirrors 1 -n lab3_split vg_summit/lab3
Logical volume lab3 converted.

[root]# lvs -o lv_name,lv_size,lv_attr,segtype,devices vg_summit/lab3
LV      LSize  Attr  Type  Devices
lab3    20.00m mwi-ao mirror lab3_mimage_0(0),lab3_mimage_1(0)

[root]# lvs -o lv_name,lv_size,lv_attr,segtype,devices vg_summit/lab3_split
LV          LSize  Attr  Type  Devices
lab3_split  20.00m -wi-ao linear /dev/vdd1(9)

[root]# lvs --segments vg_summit
LV      VG          Attr  #Str  Type  SSize
lab2    vg_summit  -wi-ao   3  striped  108.00m
lab2    vg_summit  -wi-ao   2  striped  104.00m
lab3    vg_summit  mwi-ao   2  mirror   20.00m
lab3_split vg_summit -wi-ao   1  linear   20.00m
```

lab3 has 2 images again

lab3_split

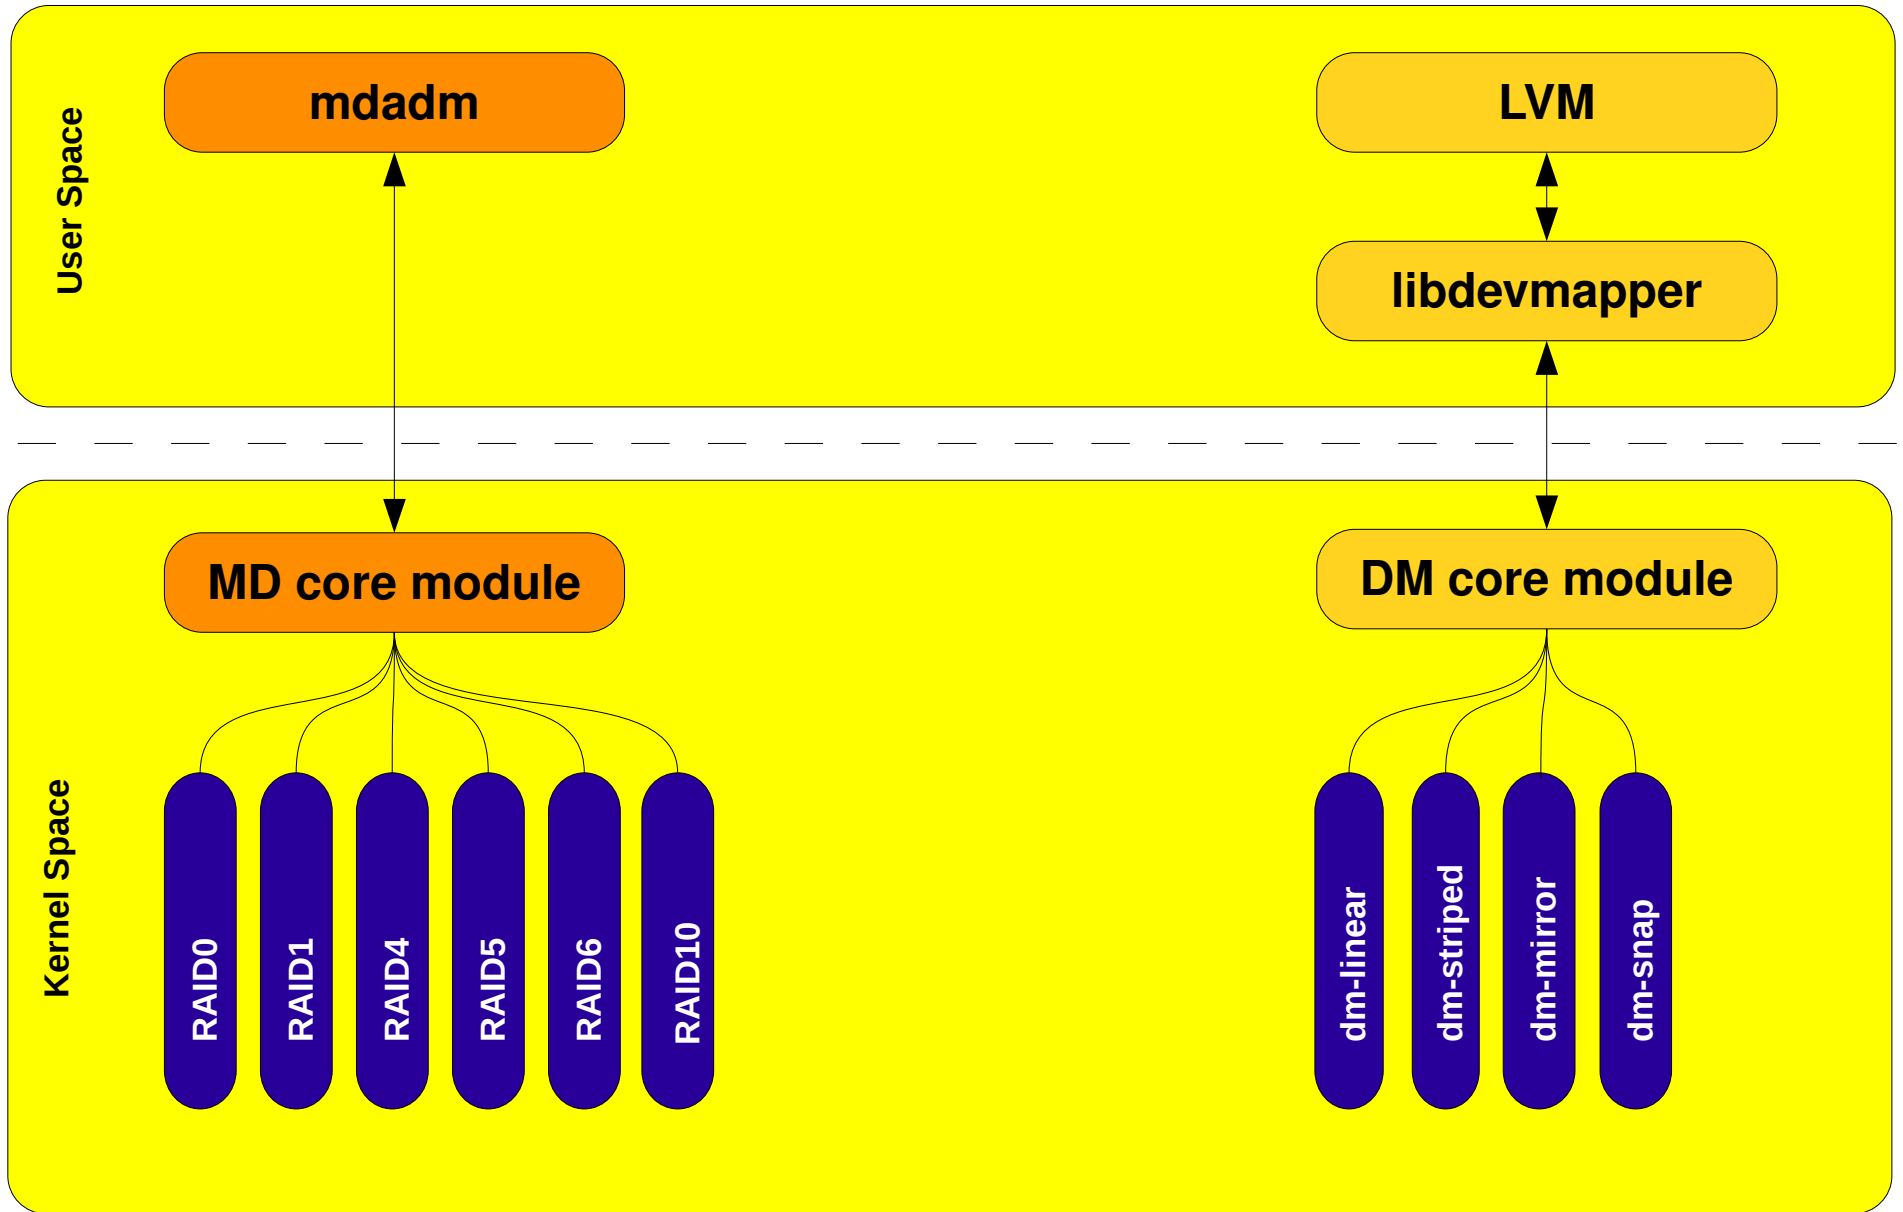
Lab 3 : Part 3 – Analyzing the filesystems

```
[root]# df
Filesystem          1K-blocks      Used Available Use% Mounted on
/dev/mapper/vg_desktop-lv_root
                    9555048    2134780    6934892   24% /
tmpfs                510404         272     510132    1% /dev/shm
/dev/vda1            495844        32891     437353    7% /boot
/dev/mapper/vg_summit-lab2
                    210350         5902     193807    3% /mnt/lab2
/dev/mapper/vg_summit-lab3
                    19827         1422     17381     8% /mnt/lab3
/dev/mapper/vg_summit-lab3_split
                    19827         1422     17381     8% /mnt/lab3_split
```

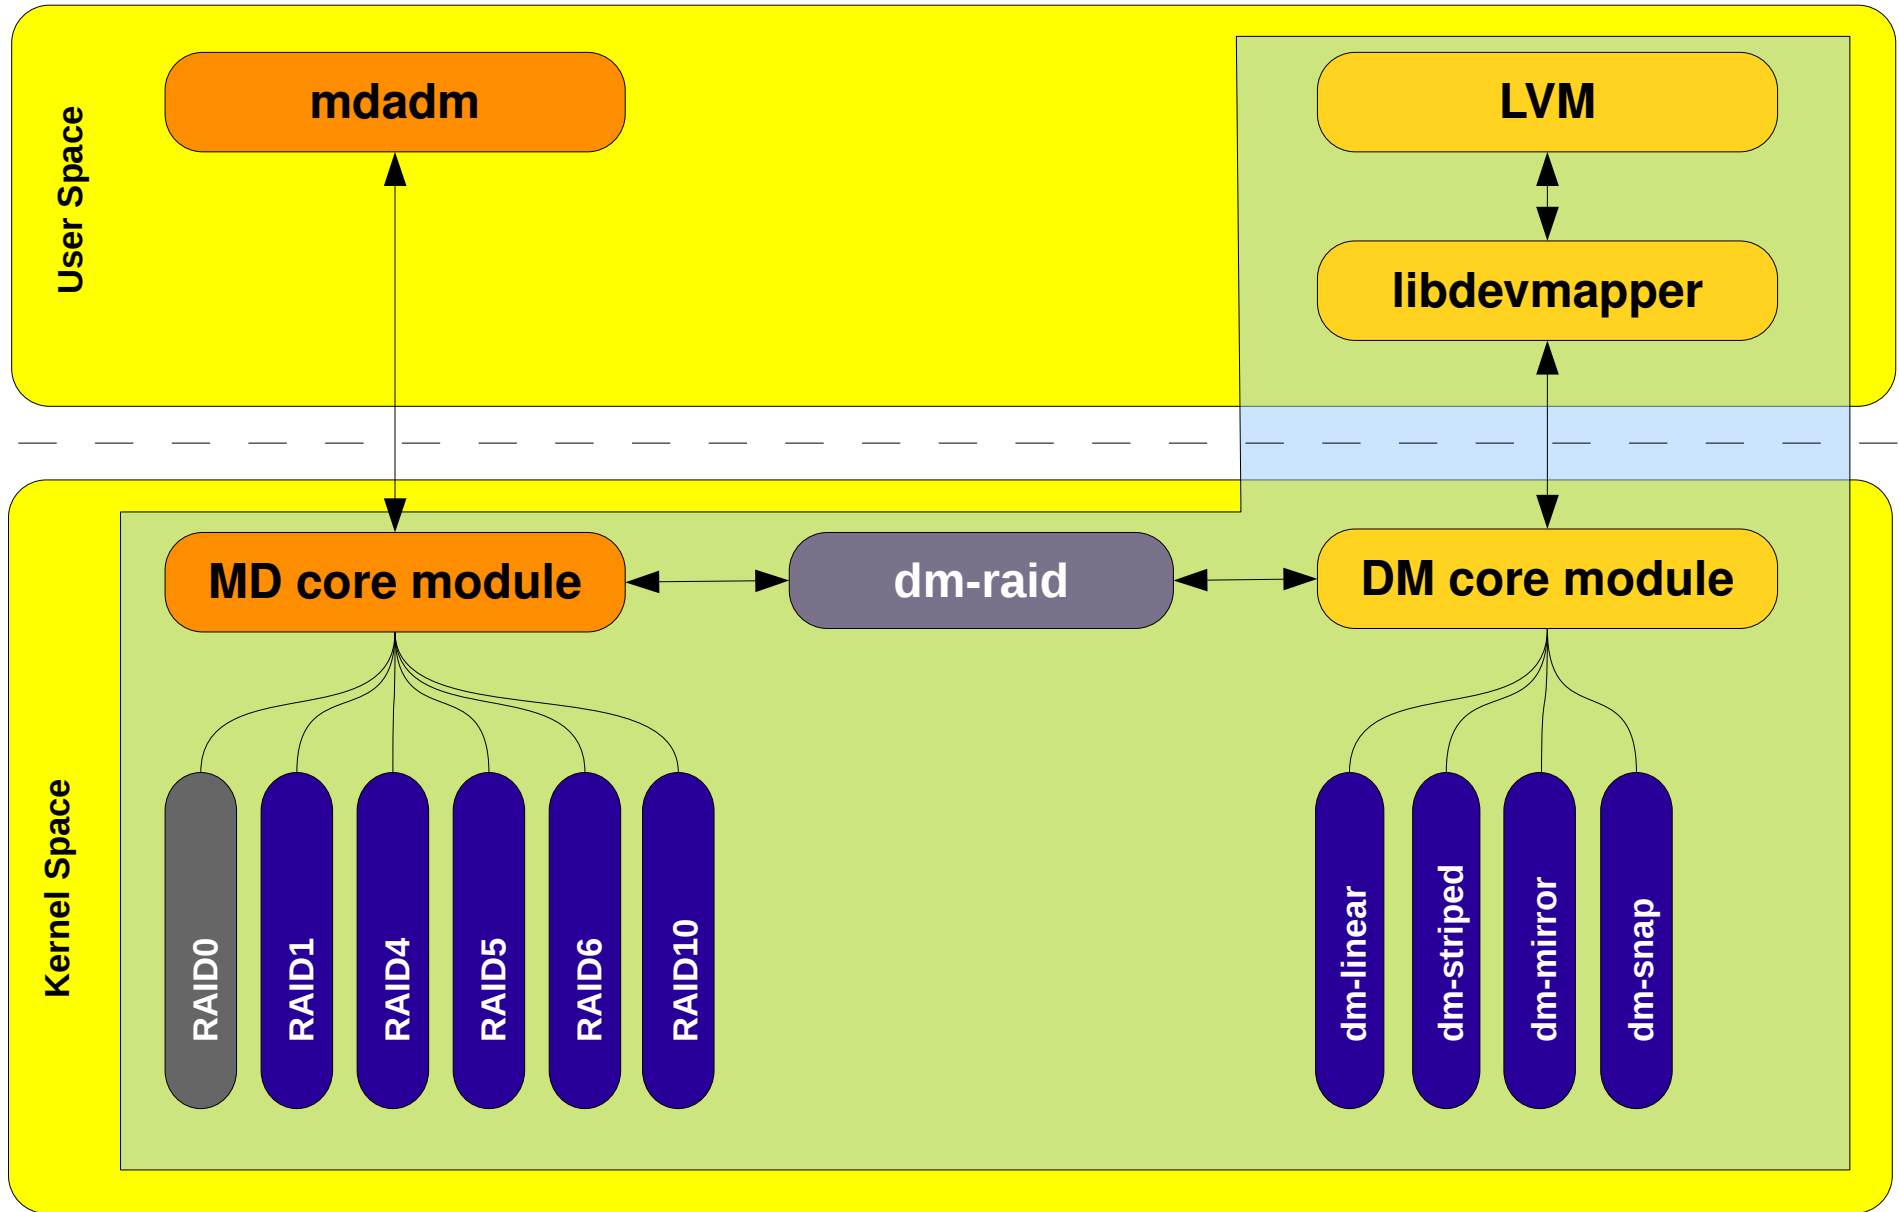
RED HAT
SUMMIT

Logical Volume Management RAID

Architectural Overview (Before)



Architectural Overview (Current)



LVM RAID 1, 4, 5, 6 & 10

- Summary:
 - RAID logical volumes provide device fault tolerance and differing I/O patterns based on the type of RAID used.
- Usage:
 - `lvcreate --type raid[456] -i <stripe_count> -n <lv_name> -L <size> <vg>`
 - `lvcreate --type raid1 -m <copy_count> -n <lv_name> -L <size> <vg>`
 - `lvcreate --type raid10 -m 1 -i <#stripes> -n <lv_name> -L <size> <vg>`
- Additional Comments:
 - '--mirrorlog' is not available for RAID logical volumes
 - RAID10 only supports 2-way mirrors (ie: -m 1)

Lab 4 – Create a RAID5 LV

Device/Partition

Physical Volume

Volume Group

Logical Volume

Filesystem

Mount

```
lvcreate --type raid5 -i 2 -n lab4 -L 20M vg_summit
```

```
mkfs -t ext4 /dev/vg_summit/lab4
```

```
mkdir -p /mnt/lab4
```

```
mount /dev/vg_summit/lab4 /mnt/lab4
```

Lab 4 – Analysis

```
[root]# lvcreate --type raid5 -i 2 -n lab4 -L 20M vg_summit
Using default stripesize 64.00 KiB
Rounding size (5 extents) up to stripe boundary size (6 extents)
Logical volume "lab4" created

[root@bp-01 ~]# lvs --segments vg_summit/lab4
LV   VG       Attr      #Str Type  SSize
lab4 vg_summit rwi-a-r-   3  raid5 24.00m
```

'i 2' and the implicit parity device make 3 total devices.

'r' attribute signifies a RAID segment type.

LVM RAID 1 – Converting a Mirror to RAID1

- Summary:
 - 'lvconvert' can be used to change the segment type from “mirror” to “raid1”
- Usage:
 - `lvconvert --type raid1 <vg>/<mirror-lv>`
- Additional Comments:
 - “raid1” provides better support for snapshots
 - “raid1” supports transient failures
 - This operation is not reversible, except via 'vgcfgrestore'

Lab 5 – Convert a Mirror to RAID1

Device/Partition

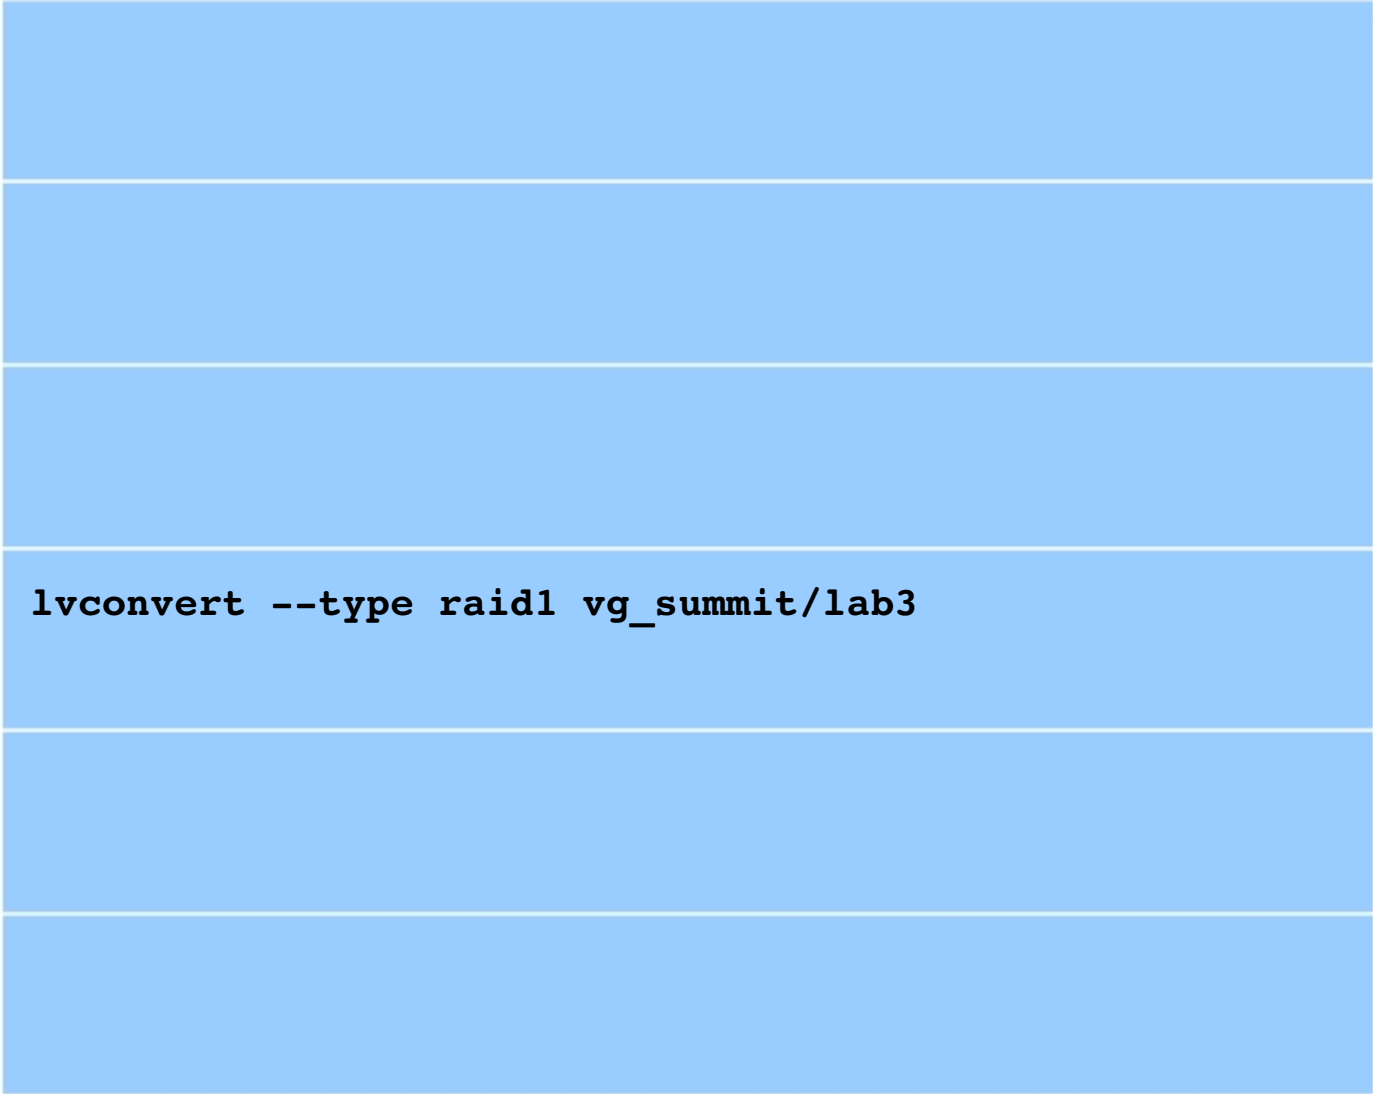
Physical Volume

Volume Group

Logical Volume

Filesystem

Mount



Lab 5 – Analyzing lvconvert

```
[root@bp-01 ~]# lvs --segments -a /dev/mapper/vg_summit-lab3*
LV          VG          Attr      #Str Type   SSize
lab3        vg_summit   mwi-a-m-   2 mirror 20.00m
[lab3_mimage_0] vg_summit   iwi-aom-   1 linear 20.00m
[lab3_mimage_1] vg_summit   iwi-aom-   1 linear 20.00m
[lab3_mlog]  vg_summit   lwi-aom-   1 linear  4.00m

[root@bp-01 ~]# lvconvert --type raid1 vg_summit/lab3

[root@bp-01 ~]# lvs --segments -a /dev/mapper/vg_summit-lab3*
LV          VG          Attr      #Str Type   SSize
lab3        vg_summit   rwi-a-m-   2 raid1  20.00m
[lab3_rimage_0] vg_summit   iwi-aor-   1 linear 20.00m
[lab3_rimage_1] vg_summit   iwi-aor-   1 linear 20.00m
[lab3_rmeta_0]  vg_summit   ewi-aor-   1 linear  4.00m
[lab3_rmeta_1]  vg_summit   ewi-aor-   1 linear  4.00m
```

“mirror” changes to “raid1”

Mirror log converts to RAID metadata areas.

Lab 5 – BONUS * Rename a Logical Volume

Device/Partition

Physical Volume

Volume Group

Logical Volume

Filesystem

Mount

```
lvrename vg_summit/lab3 vg_summit/lab5
```

```
mkdir -p /mnt/lab5
```

```
umount /mnt/lab3
```

```
mount /dev/vg_summit/lab5 /mnt/lab5
```

Lab 5 - BONUS * Analyzing the renaming of Lab3 to Lab5

```
File Edit View Search Terminal Help

[root]# lvrename vg_summit/lab3 vg_summit/lab5
Renamed "lab3" to "lab5" in volume group "vg_summit"

[root]# lvs --segments vg_summit/lab5
LV VG Attr #Str Type SSize
lab5 vg_summit rwi-aom- 2 raid1 20.00m

[root]# df -h /mnt/lab3
Filesystem Size Used Avail Use% Mounted on
/dev/mapper/vg_summit-lab3 20M 1.5M 18M 8% /mnt/lab3

[root]# umount /mnt/lab3
[root]# mkdir -p /mnt/lab5
[root]# mount /dev/vg_summit/lab5 /mnt/lab5

[root]# df -h /mnt/lab5
Filesystem Size Used Avail Use% Mounted on
/dev/mapper/vg_summit-lab5 20M 1.5M 18M 8% /mnt/lab5
```

unmount & mount to account for name change.

device & mount are updated

RED HAT
SUMMIT

Logical Volume Management

Snapshots

LVM Snapshot

- Summary:
 - Quick point-in-time copy of a filesystem without requiring complete duplication of physical storage capacity.
- Usage:
 - `lvcreate -s -L <size> -n <name> <origin-vg>/<origin-lv>`
- Additional Comments:
 - '-s' and '--snapshot' are equivalent
 - '-c/--chunksize <Size>' sets the COW copy size
(Default=4kiB)

Lab 6 – Create a Snapshot for general use

Device/Partition

Physical Volume

Volume Group

Logical Volume

Filesystem

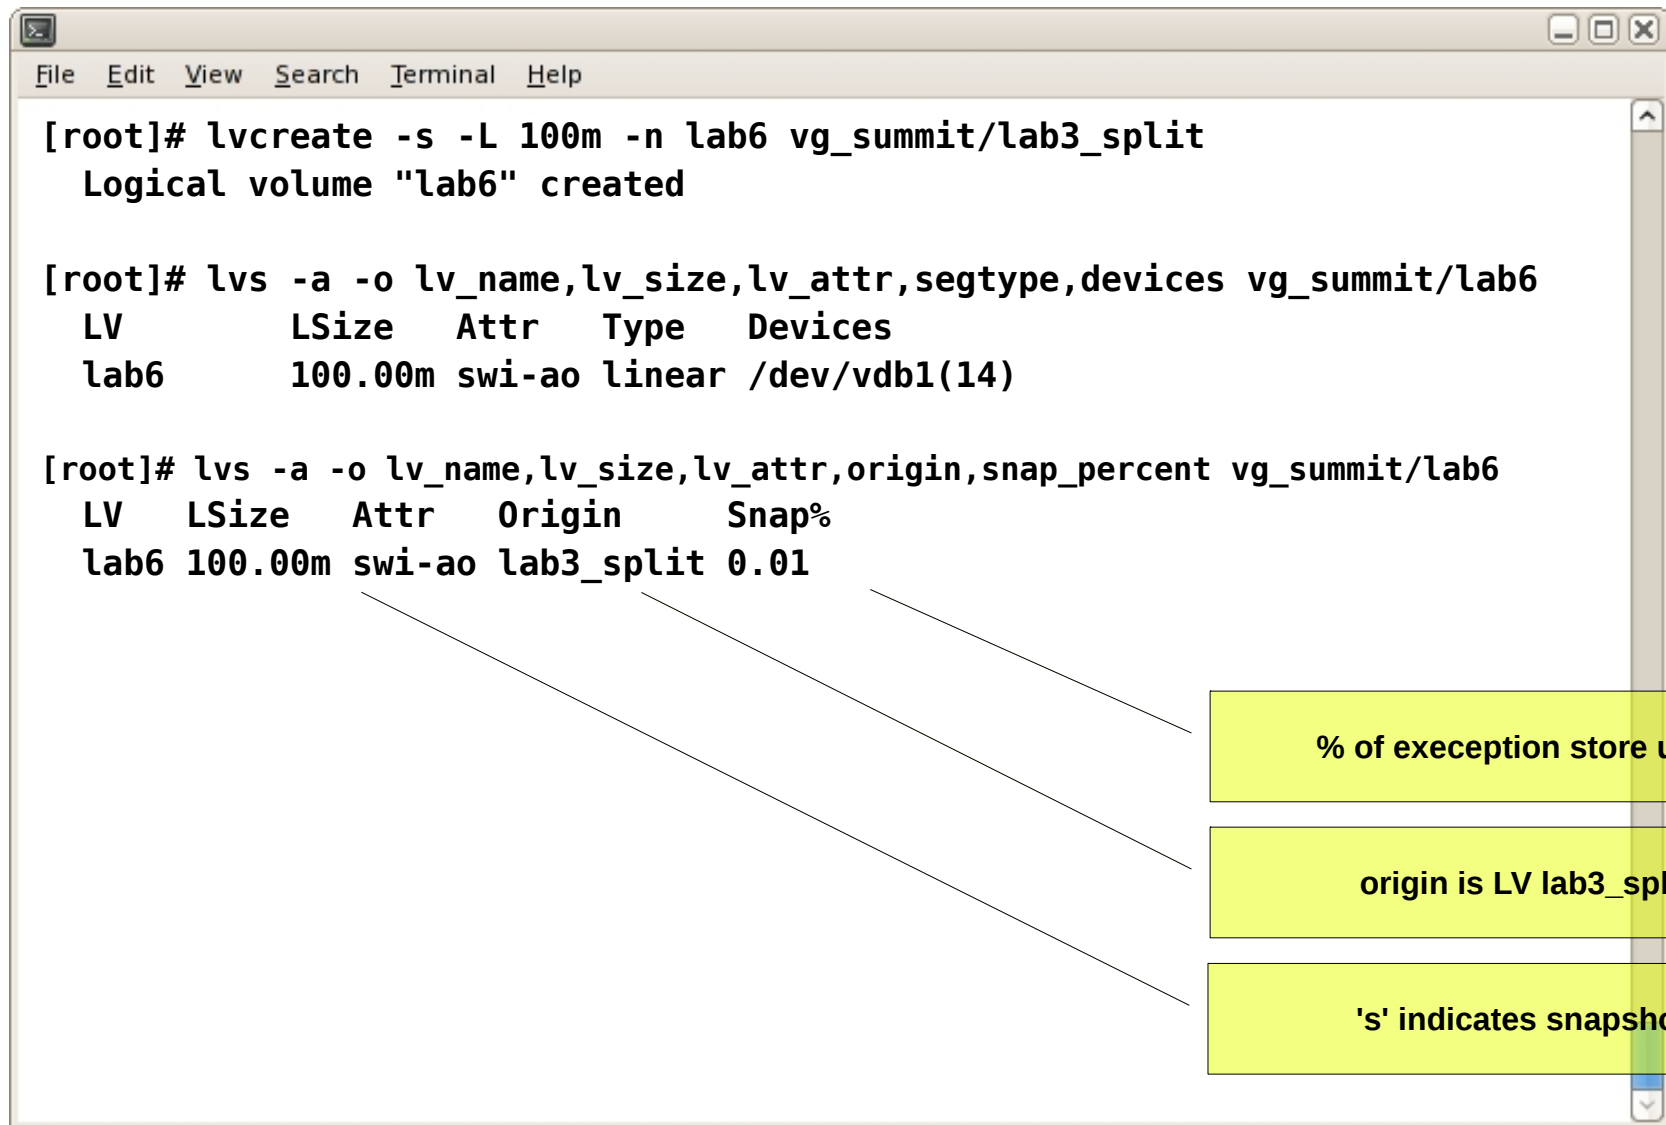
Mount

```
lvcreate -s -L 100m -n lab6 vg_summit/lab3_split
```

```
mkdir -p /mnt/lab6
```

```
mount /dev/vg_summit/lab6 /mnt/lab6
```

Lab 6 – Analyzing lvcreate



A terminal window showing the execution of the `lvcreate` command and subsequent `lvs` commands. The output shows the creation of a logical volume named `lab6` with a size of 100m and attributes `swi-ao`. The `lvs` command shows the origin of the snapshot as `lab3_split` and the snapshot percentage as `0.01`. Three callout boxes on the right explain the output: the `0.01` indicates the percentage of exception store used, the `lab3_split` origin indicates the source LV, and the `s` in the attributes indicates a snapshot.

```
[root]# lvcreate -s -L 100m -n lab6 vg_summit/lab3_split
Logical volume "lab6" created

[root]# lvs -a -o lv_name,lv_size,lv_attr,segtype,devices vg_summit/lab6
LV          LSize  Attr  Type  Devices
lab6       100.00m swi-ao linear /dev/vdb1(14)

[root]# lvs -a -o lv_name,lv_size,lv_attr,origin,snap_percent vg_summit/lab6
LV  LSize  Attr  Origin    Snap%
lab6 100.00m swi-ao lab3_split 0.01
```

% of exception store used

origin is LV lab3_split

's' indicates snapshot

Lab 6 : Remove some content from origin volume

```
ls -god /mnt/lab3_split/sysconfig/iptables*  
rm -f /mnt/lab3_split/sysconfig/iptables*  
ls -god /mnt/lab3_split/sysconfig/iptables*  
ls -god /mnt/lab6/sysconfig/iptables*
```

Lab 6 – Analyzing the filesystems

```
File Edit View Search Terminal Help

[root]# ls -god /mnt/lab3_split/sysconfig/iptables*
-rw----- . 1 476 Jun 23 23:02 /mnt/lab3_split/sysconfig/iptables
-rw----- . 1 1740 Jun 23 23:02 /mnt/lab3_split/sysconfig/iptables-config
-rw----- . 1 476 Jun 23 23:02 /mnt/lab3_split/sysconfig/iptables.old

[root]# rm -f /mnt/lab3_split/sysconfig/iptables*

[root]# ls -god /mnt/lab3_split/sysconfig/iptables*
ls: cannot access /mnt/lab3_split/sysconfig/iptables*: No such file or directory

[root]# ls -god /mnt/lab6/sysconfig/iptables*
-rw----- . 1 476 Jun 23 23:02 /mnt/lab6/sysconfig/iptables
-rw----- . 1 1740 Jun 23 23:02 /mnt/lab6/sysconfig/iptables-config
-rw----- . 1 476 Jun 23 23:02 /mnt/lab6/sysconfig/iptables.old

[root]# lvs -a -o lv_name,lv_size,lv_attr,origin,snap_percent vg_summit/lab6
LV   LSize  Attr   Origin Snap%
lab6 100.00m swi-ao lab3   0.06
```

Files still available in snapshot

Notice some space is now used

LVM Snapshot / Merge

- Summary:
 - Merging a snapshot back to its origin provides a convenient method to return a filesystem to its previous state.
- Usage:
 - `lvconvert --merge <vg>/<snap-lv>`
- Additional Comments:
 - Merging requires LV deactivation & reactivation
 - Merging the root filesystem requires a reboot

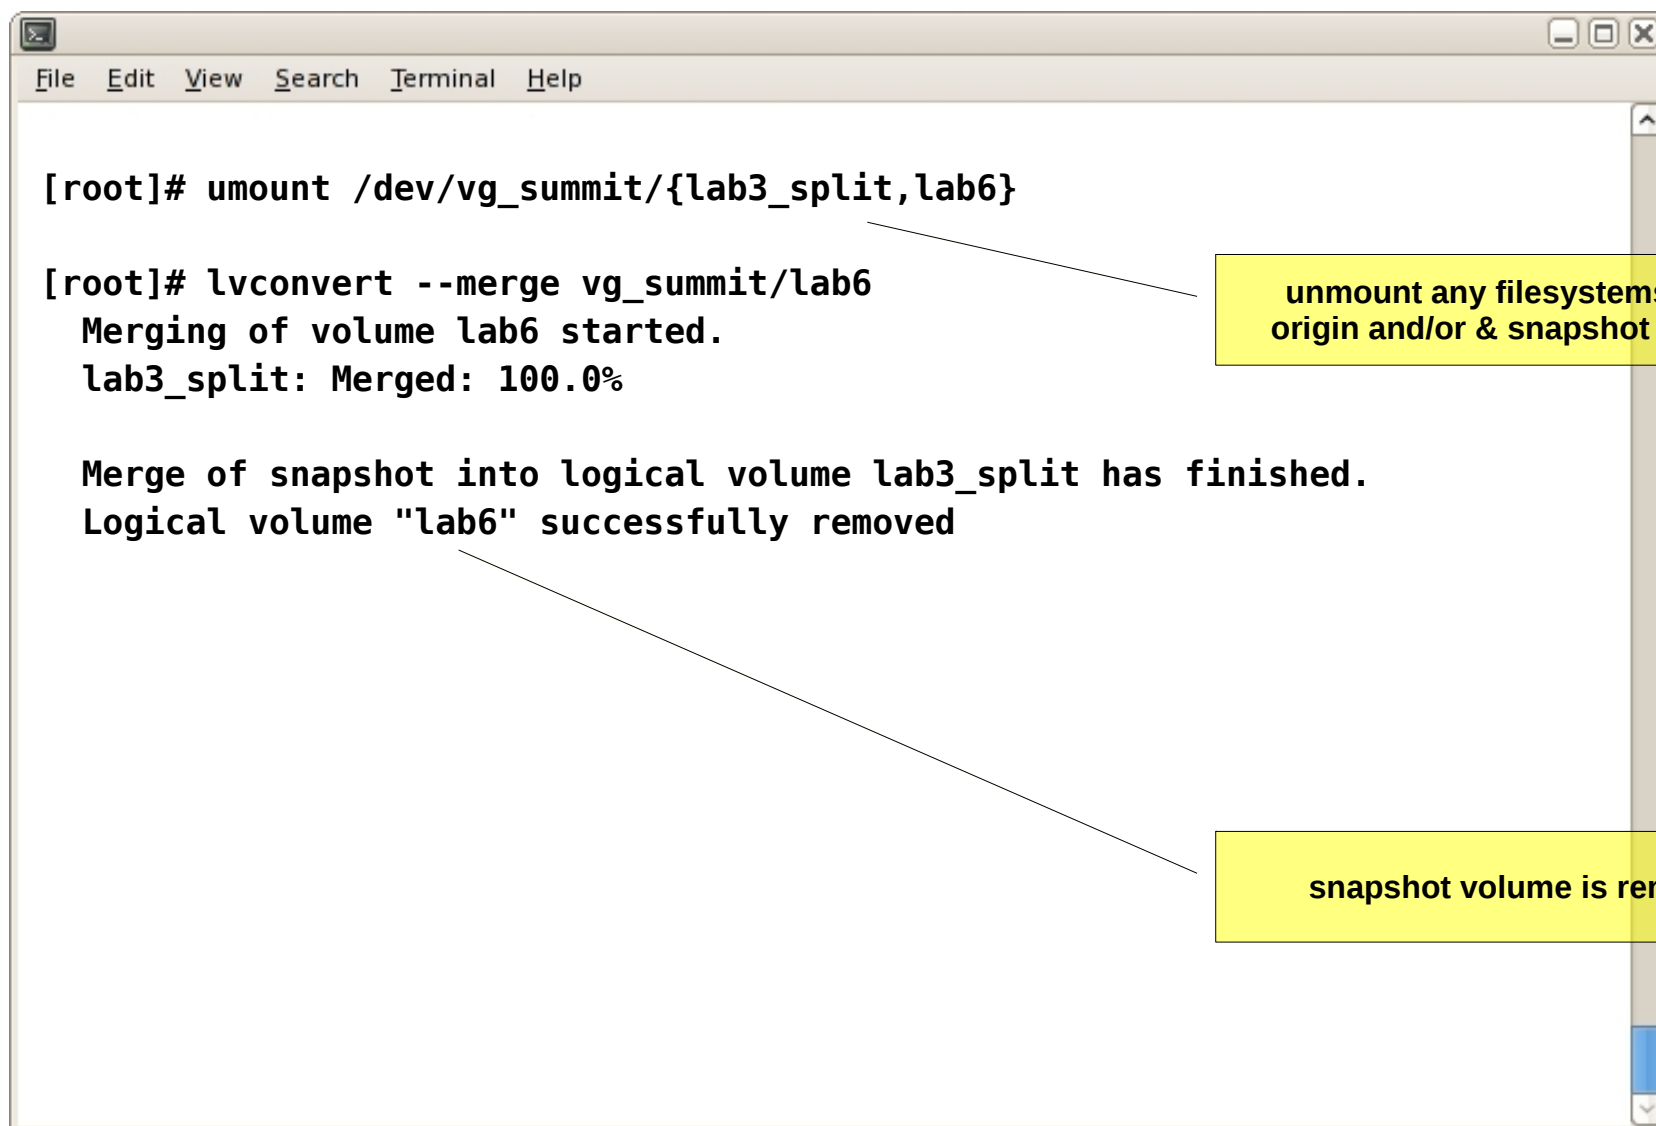
Lab 7 : Merge LV to Previous State

```
umount /dev/vg_summit/{lab3_split,lab6}
```

```
lvconvert --merge vg_summit/lab6
```

```
mount /dev/vg_summit/lab3_split /mnt/lab3_split
```

Lab 7 – Analyzing Merge



```
[root]# umount /dev/vg_summit/{lab3_split,lab6}

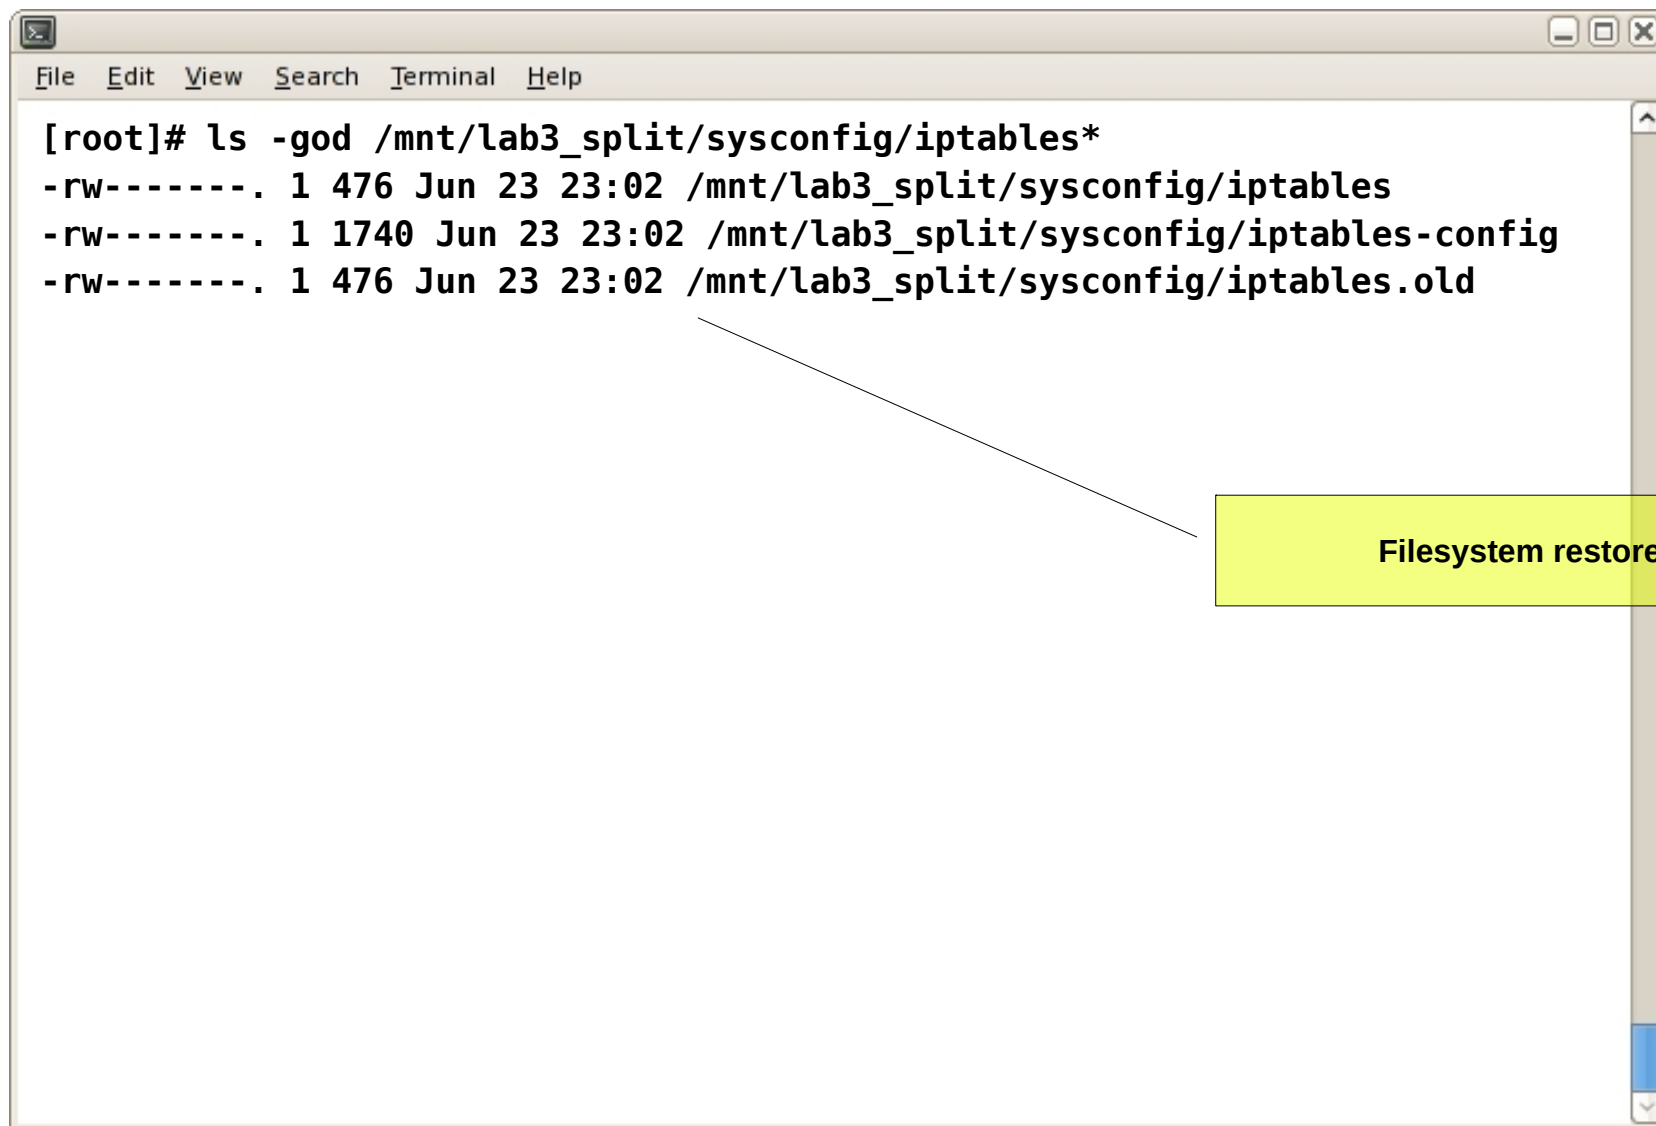
[root]# lvconvert --merge vg_summit/lab6
Merging of volume lab6 started.
lab3_split: Merged: 100.0%

Merge of snapshot into logical volume lab3_split has finished.
Logical volume "lab6" successfully removed
```

umount any filesystems using origin and/or & snapshot volumes

snapshot volume is removed

Lab 7 – Analyzing the filesystems



```
[root]# ls -god /mnt/lab3_split/sysconfig/iptables*  
-rw-----. 1 476 Jun 23 23:02 /mnt/lab3_split/sysconfig/iptables  
-rw-----. 1 1740 Jun 23 23:02 /mnt/lab3_split/sysconfig/iptables-config  
-rw-----. 1 476 Jun 23 23:02 /mnt/lab3_split/sysconfig/iptables.old
```

The terminal window displays the output of the command `ls -god /mnt/lab3_split/sysconfig/iptables*`. The output shows three files: `/mnt/lab3_split/sysconfig/iptables`, `/mnt/lab3_split/sysconfig/iptables-config`, and `/mnt/lab3_split/sysconfig/iptables.old`. A yellow callout box with the text "Filesystem restored" is connected to the terminal output by a thin black line.

LVM Sparse Volumes

- Summary:
 - Snapshots can be used as read/write volumes and can be created without an origin. This provides a way to create large devices without the need for the full physical space.
- Usage:
 - `lvcreate -s -V <virtual_size> -L <real_size> <VG>`
- Additional Comments:
 - A virtual 'zero' device is used as the origin for sparse volumes
 - When the real space is used-up, the volume is dead.

**** WARNING ** 100% capacity = 100% data loss**

Lab 8 – Create a sparse logical volume

Device/Partition

Physical Volume

Volume Group

Logical Volume

Filesystem

Mount

```
lvcreate -s -V 1G -L 100M -n lab8 vg_summit
```

```
mkfs -t ext4 /dev/vg_summit/lab8
```

```
mkdir -p /mnt/lab8
```

```
mount /dev/vg_summit/lab8 /mnt/lab8
```

Lab 7 – Analyzing creation of a sparse logical volume

```
File Edit View Search Terminal Help

[root]# lvcreate -s -V 1G -L100M -n lab8 vg_summit
Logical volume "lab8" created.

[root]# mkfs -t ext4 /dev/vg_summit/lab8 > /dev/null

[root]# mkdir -p /mnt/lab8

[root]# mount /dev/vg_summit/lab8 /mnt/lab8

[root]# df -h /mnt/lab8
Filesystem                Size      Used Avail Use% Mounted on
/dev/mapper/vg_summit-lab8 1008M    34M   924M   4% /mnt/lab8

[root]# lvs -o name,vg_name,attr,size,origin,snap_percent vg_summit/lab8
LV   VG      Attr      LSize  Origin      Snap%
lab8 vg_summit swi-aos-- 100.00m [lab7_vorigin] 48.59
```

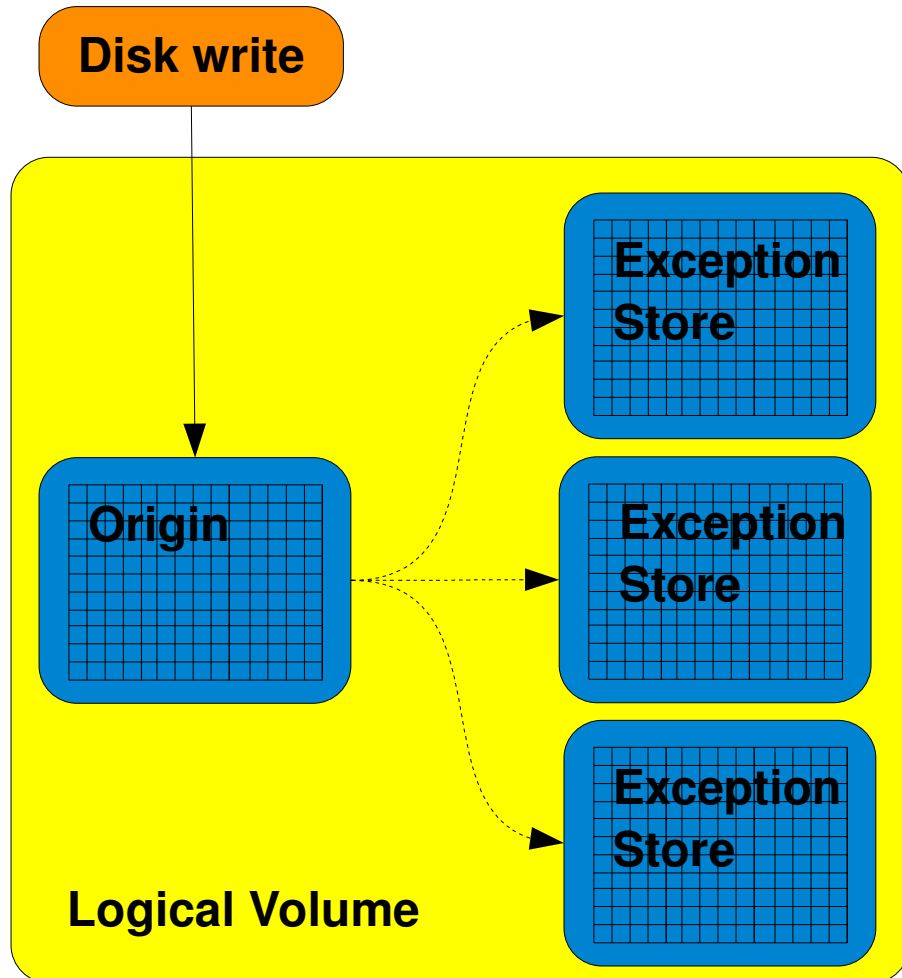
RED HAT
SUMMIT

Logical Volume Management

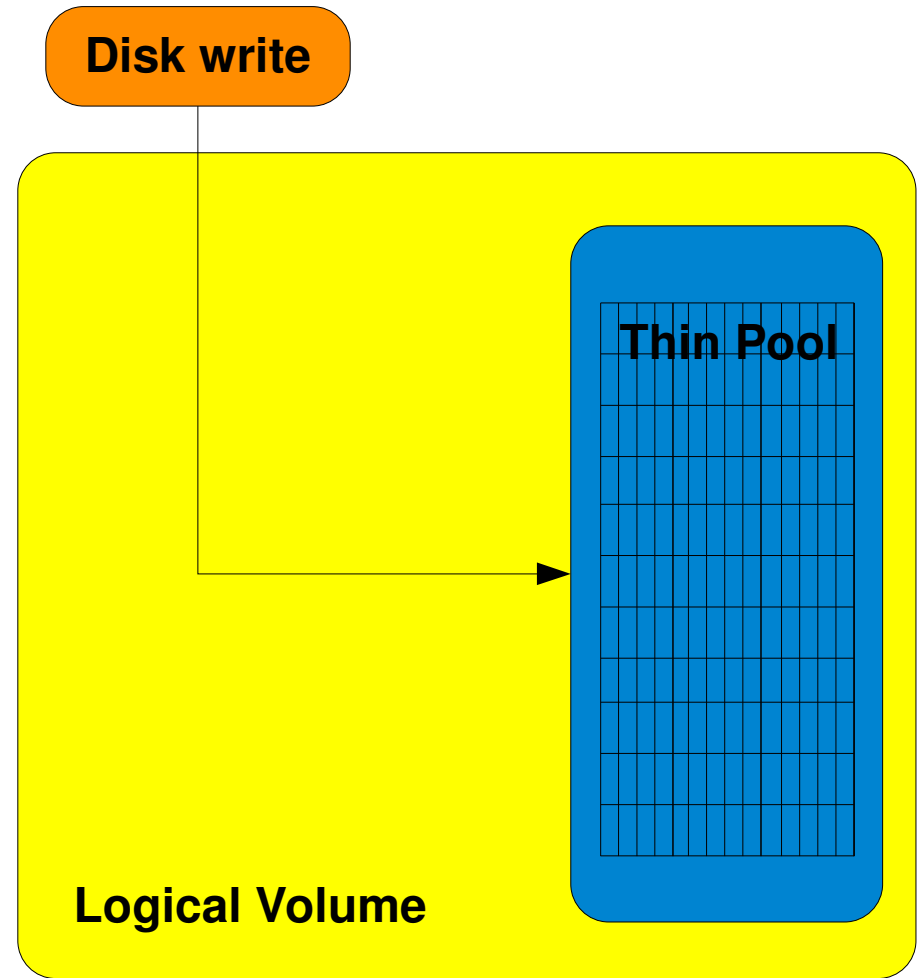
Thin Provisioning

Architectural Overview

Classic Snapshot



Thin Snaps



LVM Thin Logical Volumes and Snapshots

- Summary:
 - Thin logical volumes provide the means to reserve capacity upfront & consume necessary physical space on demand.
- Usage:
 - `lvcreate -T <vg>/<pool_name> -L <size> -V <size> -n <lv_name>`
 - `lvcreate -s <vg>/<origin_lv> -n <thin_snap_name>`
- Additional Comments:
 - No size specification for thin snapshots
 - Thin snaps share a common storage pool which makes it very efficient for large numbers of snapshots

Lab 9: Part 1 – Create a Thin Logical Volume

Device/Partition

Physical Volume

Volume Group

Logical Volume

Filesystem

Mount

```
lvcreate -T vg_summit/lab9_pool -L 100M -V 1G -n lab9
```

```
mkfs.ext4 /dev/vg_summit/lab9
```

```
mkdir -p /mnt/lab9
```

```
mount /dev/vg_summit/lab9 /mnt/lab9
```

Lab 9: Part 1 – Analyzing creation of a thin logical volume

```
File Edit View Search Terminal Help

[root]# lvcreate -T vg_summit/lab9_pool -L 100M -V 1G -n lab9
Logical volume "lab9" created

[root]# mkfs.ext4 /dev/vg_summit/lab9 > /dev/null

[root]# mkdir -p /mnt/lab9

[root]# mount /dev/vg_summit/lab9 /mnt/lab9

[root]# df -h /mnt/lab9
Filesystem                Size      Used Avail Use% Mounted on
/dev/mapper/vg_summit-lab9 1008M    34M   924M   4% /mnt/lab8

[root]# lvs -o name,vg_name,attr,size,pool_lv,data_percent vg_summit
LV          VG          Attr      LSize   Pool        Data%
lab2        vg_summit   -wi-ao--- 40.00m
lab3_split vg_summit   -wi-ao--- 20.00m
lab4        vg_summit   rwi-aor-- 24.00m
lab5        vg_summit   rwi-aor-- 20.00m
lab8        vg_summit   swi-aos-- 100.00m                48.59
lab9        vg_summit   Vwi-aotz-  1.00g   lab9_pool    4.76
lab9_pool   vg_summit   twi-a-tz- 100.00m                48.75
```

Lab 9: Part 2 – Create a Thin Snapshot

Device/Partition

Physical Volume

Volume Group

Logical Volume

Filesystem

Mount

```
lvcreate -s vg_summit/lab9 -n lab9_snap
```

```
mkdir -p /mnt/lab9_snap
```

```
mount /dev/vg_summit/lab9_snap /mnt/lab9_snap
```

Lab 9: Part 2 – Analyzing creation of a thin logical volume

```
[root]# lvcreate -s vg_summit/lab9 -n lab9_snap
Logical volume "lab9_snap" created.

[root]# mkdir -p /mnt/lab9_snap

[root]# mount /dev/vg_summit/lab9_snap /mnt/lab9_snap

[root]# df -h /mnt/lab9_snap
Filesystem                Size      Used Avail Use% Mounted on
/dev/mapper/vg_summit-lab9_snap 1008M    34M   924M   4% /mnt/lab9_snap

[root]# lvs -o name,vg_name,attr,size,pool_lv,origin,data_percent vg_summit
LV          VG          Attr      LSize   Pool        Origin      Data%
lab2        vg_summit   -wi-ao--- 40.00m
lab3_split  vg_summit   -wi-ao--- 20.00m
lab4        vg_summit   rwi-aor-- 24.00m
lab5        vg_summit   rwi-aor-- 20.00m
lab8        vg_summit   swi-aos-- 100.00m          [lab8_vorigin] 48.59
lab9        vg_summit   Vwi-aotz-  1.00g  lab9_pool
lab9_pool   vg_summit   twi-a-tz- 100.00m          48.81
lab9_snap   vg_summit   Vwi-aotz-  1.00g  lab9_pool  lab9        4.76
```

RED HAT
SUMMIT

Logical Volume Management Wrap Up

What we DID NOT talk about

- Failed device recovery procedures
- `vgcfgbackup` & `vgcfgrestore`
- HA LVM
- Cluster LVM

Documentation

- Online User Guides
 - <http://docs.redhat.com/docs/en-US/index.html>
 - Logical_Volume_Manager_Administration
 - Storage_Administration_Guide
- Customer Portal
 - <https://access.redhat.com/home>

Additional Self Help

- Man & Info Pages

`man <topic>`

`man -k topic`

`man -s <section> <topic>`

`info <topic>`

- Installed Docs

`/usr/share/doc/*`

Be sure to install kernel-doc rpm package

RPMs include docs, man pages, etc...

`rpm -qil <package-name> | more`

**RED HAT
SUMMIT**

**Please Fill Out Your Survey
&
Thank You!**

**RED HAT
SUMMIT**

Additional Exercises #1

Using snapshot & merge to roll-back system updates

AdLab 1 : Part 1 – Cleanup all previous labs

```
umount /mnt/*  
vgremove -f vg_summit
```

AdLab 1 : Part 2 – Prepare for OS volume snapshot

```
vgextend vg_desktop /dev/vd{b,c,d,e,f}1
```

```
lvcreate -L 100M -n adlab1_arc vg_desktop
```

```
mkfs -t ext4 /dev/vg_desktop/adlab1_arc
```

```
mkdir -p /mnt/adlab1_arc
```

```
mount /dev/vg_desktop/adlab1_backup /mnt/adlab1_arc
```

Add the PVs to vg_desktop
NOT vg_summit !!!

A volume for a filesystem

AdLab 1 : Part 3 – Backup /boot

```
umount /boot
```

```
dd if=/dev/vda1 | gzip -c9 - > /mnt/adlab1_arc/boot.dd.gz
```

```
mount /boot
```

```
lvcreate -s -L 500M -n adlab1_snap vg_desktop/lv_root
```

Back up /boot

Create snapshot of root volume

AdLab 1 : Part 3 – Analyzing the system

```
[root]# lvs -a -o lv_name,lv_size,lv_attr,origin,snap_percent vg_desktop
LV          LSize  Attr      Origin  Snap%
adlab1_arc  100.00m -wi-ao--
adlab1_snap 500.00m swi-a-s- lv_root  0.03
lv_root     9.26g  owi-aos-
lv_swap     256.00m -wi-ao--

[root]# ls -god /usr/share/doc/yum*
drwxr-xr-x. 2 4096 Jun 19 22:44 /usr/share/doc/yum-3.2.29
drwxr-xr-x. 2 4096 Jun 19 22:44 /usr/share/doc/yum-metadata-parser-1.1.2
drwxr-xr-x. 2 4096 Jun 19 22:47 /usr/share/doc/yum-plugin-security-1.1.30
drwxr-xr-x. 2 4096 Jun 19 22:45 /usr/share/doc/yum-rhn-plugin-0.9.1
drwxr-xr-x. 2 4096 Jun 19 22:47 /usr/share/doc/yum-utils-1.1.30

[root]# tail -1 /boot/grub/grub.conf
        initrd /initramfs-2.6.32-220.el6.x86_64.img

[root]# cat /etc/redhat-release
Red Hat Enterprise Linux Server release 6.3 (Santiago)
```

snapshot of lv_root volume

AdLab 1 : Part 4 – Apply some changes / updates

```
echo "##LAB7 COMMENT##" >> /boot/grub/grub.conf
```

```
echo "##LAB7 COMMENT##" > /etc/redhat-release
```

```
rm -rf /usr/share/doc/yum*
```

AdLab 1 : Part 4 – Analyzing the system after updates

```
[root]# lvs -a -o lv_name,lv_size,lv_attr,origin,snap_percent vg_desktop
LV          LSize  Attr    Origin  Snap%
lab7_backup 100.00m -wi-ao--
lab7_snap   500.00m swi-a-s- lv_root  0.08
lv_root     9.26g  owi-aos-
lv_swap     256.00m -wi-ao--
```

[root]# ls -god /usr/share/doc/yum*
ls: cannot access /usr/share/doc/yum*: No such file or directory

```
[root]# tail -1 /boot/grub/grub.conf
"##LAB7 COMMENT##"
```

```
[root]# cat /etc/redhat-release
"##LAB7 COMMENT##"
```

Callouts from the terminal to the right:

- Line 4 of the `lvs` output: snap size has increased
- Output of `ls`: files deleted
- Output of `tail`: file in /boot modified
- Output of `cat`: file in / modified

AdLab 1 : Part 5 – Merge snapshot to origin

```
umount /boot
```

```
gzip -dc /mnt/lab7_arc/boot.dd.gz | dd of=/dev/vda1
```

```
fsck /dev/vda1
```

```
mount /boot
```

```
lvconvert --merge vg_desktop/adlab1_snap
```

```
reboot
```

AdLab 1 : Part 5 – Analyzing the system post reboot

```
File Edit View Search Terminal Help

[root]# lvs -a -o lv_name,lv_size,lv_attr,origin,snap_percent vg_desktop
LV          LSize  Attr   Origin Snap%
adlab1_arc  100.00m -wi-a---
lv_root     9.26g  -wi-ao--
lv_swap     256.00m -wi-ao--

[root]# ls -god /usr/share/doc/yum*
drwxr-xr-x. 2 4096 Jun 25 00:35 /usr/share/doc/yum-3.2.29
drwxr-xr-x. 2 4096 Jun 19 22:44 /usr/share/doc/yum-metadata-parser-1.1.2
drwxr-xr-x. 2 4096 Jun 25 00:41 /usr/share/doc/yum-plugin-security-1.1.30
drwxr-xr-x. 2 4096 Jun 25 00:36 /usr/share/doc/yum-rhn-plugin-0.9.1
drwxr-xr-x. 2 4096 Jun 25 00:41 /usr/share/doc/yum-utils-1.1.30

[root]# tail -1 /boot/grub/grub.conf
initrd /initramfs-2.6.32-220.el6.x86_64.img

[root]# cat /etc/redhat-release
Red Hat Enterprise Linux Server release 6.4 (Santiago)
```

lab7_snap snapshot is gone

files returned

/boot restored