


# **Manage your disk space... for free :)**

*Julien Wallior*

**Plug Central**  
**11.1.2006**

# Agenda




- Background
- RAID
  - Some basics
  - Practice
  - Booting on a raid device
- LVM
  - What is that?
  - How it works
- Hardware raid... if you really want it

# Background



- Just got out of college
  - I don't want a pay for anything
  - I don't need that much performance
  - If it breaks... it's my problem
- Does it mean this presentation is useless?
  - Of course not...
  - I've been using these technologies for two years without issue

# *RAID: The Basics*




- Redundant Array of Inexpensive Disks
- Two usage:
  - Performance
  - Redundancy
- Different RAID level depending on your needs
  - Tradeoffs

# Basic Levels




- RAID 0 (stripping)

- Performance ++
- Redundancy --




- RAID 1 (mirroring)

- Performance
  - Read ++
  - Write --
- Redundancy ++






# Fancier (or Funkier)



- RAID 5

- Distributed Parity

- Performance
  - Read ++
  - Write -
- Redundancy +



- Less popular

- Dedicated parity disk

- RAID 3

- Byte level stripping

- RAID 4


- Block level stripping

- Two parity disks

- RAID 6

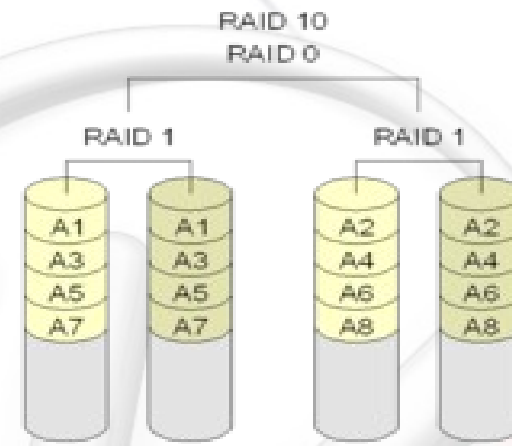
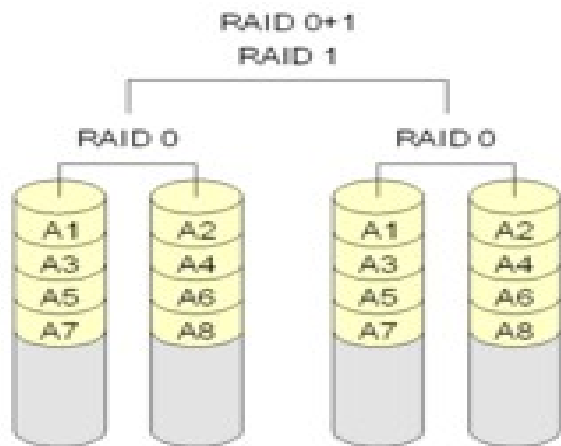
- Like RAID 5

# The Combos



- RAID 0+1

- RAID 1+0



What's the difference???


# *mdadm...*



- **Some nice features:**
  - integrated in the OS
  - spare drives
  - email when a disk fails
  - periodic rebuild
- **Enough theory for now...**



# Let's play a little bit



- Create a raid 5 array with one hot spare
  - `mdadm --create /dev/md0 --raid-devices=3 --spare-device=1 --level=5 /dev/ram0 /dev/ram1 /dev/ram2 /dev/ram3`
- Check the state of the array
  - `mdadm --detail /dev/md0`
  - `cat /proc/mdstat`
- Simulate a drive failure
  - `mdadm --manage /dev/md0 --set-faulty /dev/ram0`

# *Let's play a little bit*




- Hot remove a drive
  - `mdadm --manage /dev/md0 --remove /dev/ram0`
- Hot add a drive
  - `mdadm --manage /dev/md0 --add /dev/ram4`
- Some performance measurements
  - Raid 5 vs Raid 1 write: 7 times slower !

# Booting on it?



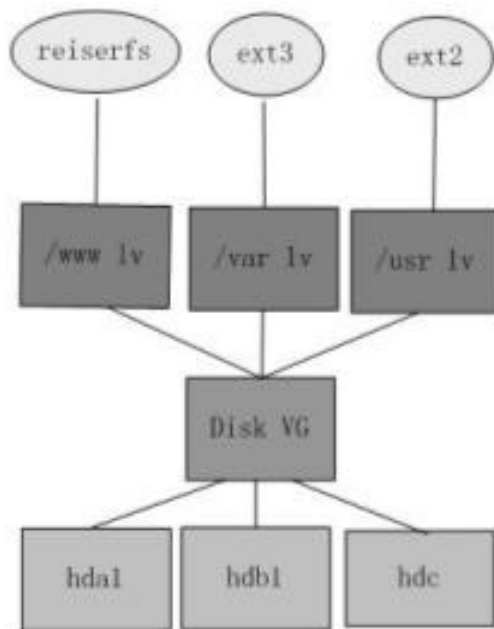
- It works...
- Little tricky grub setup (mbr not mirrored)
  - grub> device (hd0) /dev/hda
  - grub> root (hd0,0)
  - grub> setup (hd0)
  - grub> device (hd0) /dev/hdb
  - grub> root (hd0,0)
  - grub> setup (hd0)
  - grub> exit

# *LVM: what is that?*



- Logical Volume Manager
- LVM2 is standard in kernel 2.6

- Physical volume (pv...)
- Volume group (vg...)
- Logical volume (lv...)




# *Why would I use it?*



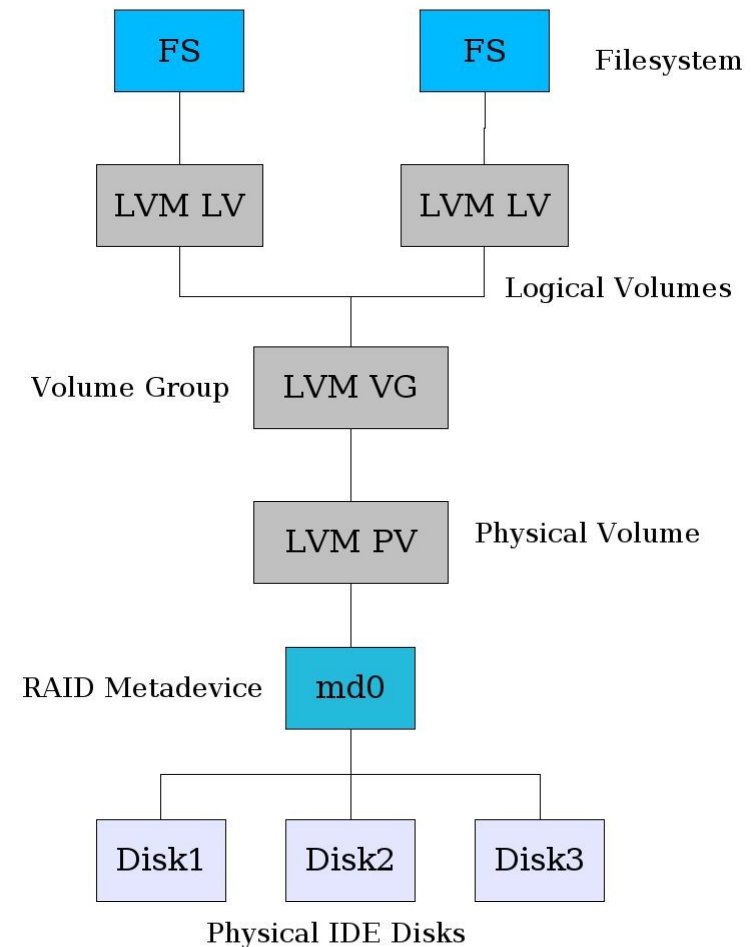
- Partition spread across multiple disks
  - different than RAID0 (no striping)
- Resize partition of the fly
  - you just need to unmount it
  - don't forget to resize the filesystem (resize2fs)
- Snapshots
  - very useful for backups
  - requires free space in the volume group




# Let's see something



- Enough talking...
  - Our goal:
    - RAID for redundancy
    - LVM for flexibility
- Demo...



# *Physical Volume*



- Creation
  - `pvcreate /dev/md0`
- Get config
  - `pvs`



# Volume group



- **Creation**
  - `vgcreate vg1 /dev/md0`
- **Get config**
  - `vgs`
- **Add a new physical volume**
  - `vgextend vg1 /dev/ram5`

# *Logical volume*




- **Creation**
  - `lvcreate --name lv1 --size 80M vg1`
  - `mkfs.ext3 /dev/mapper/vg1-lv1`
  - `mount /dev/mapper/vg1-lv1 /mnt`
- **Get state**
  - `lvs`
- **Resize**
  - `resize2fs /dev/mapper/vg1-lv1 80M` (don't forget it)
  - `lvresize -L -20M /dev/vg1/lv1`

# Snapshot




- Takes a snapshot of a logical volume
  - modprobe dm-snapshot
- Creation
  - `lvcreate --snapshot --name lv1-snapshot1 --size 80M /dev/vg1/lv1`
- Deletion
  - `lvremove /dev/vg1/lv1-snapshot1`





- Better performances... more expensive
- On HP Proliant servers
- Utility hpacucli for linux
  - `ctrl slot=0 pd all show`
  - `ctrl slot=0 ld all show`
  - `ctrl slot=0 create type=ld drives=2:0,2:1 raid=1`
- Hot plug and hot swap the disks...
- Basically it works and it's no fun at all :)

# Conclusion



- Personal experience:
  - You can really do a lot
  - Excellent mean to learn how it works
  - It's amazing how well it works
- Sorry for the french accent :)
- One last thing: make a backup before messing around with these things.

*Thank you*



**Questions ?**