

Linux and RAID


Doug Warner
doug@warner.fm
<http://doug.warner.fm>
Barracuda Networks

RAID Overview

RAID IS NOT BACKUP


RAID Types

RAID-0




http://en.wikipedia.org/wiki/File:RAID_0.svg

RAID-1




http://en.wikipedia.org/wiki/File:RAID_1.svg

RAID-5



http://en.wikipedia.org/wiki/File:RAID_5.svg

RAID-6



http://en.wikipedia.org/wiki/File:RAID_6.svg

Create Some RAID

- Partition your drives (fdisk, parted)
- **mdadm** util
- Format your new partition!
- Mount your partition!

Create Some RAID - mdadm

- `mdadm --create -help`
 - `/dev/mdX`
 - `--bitmap=internal`
 - `--chunk=256`
 - `--level=6`
 - `--raid-devices=6`
 - `/dev/sdX /dev/sdY ...`

Check Out Your RAID

- **/proc/mdstat**
- **/sys/block/mdX/md/**
 - sync_speed
 - mismatch_cnt
 - sync_action

Recovering From Drive Failures



<http://wearehugh.com/public/2008/05/dont-panic.jpg>

Recovering From Drive Failures

Your Data Lives On

- Write Cache
 - Battery Back-Up (BBU)?
 - Write Cache + nobarrier
 - No BBU?
 - Write-Through + barriers

Recovering From Drive Failures [Faster]

Check Your Data

Your Data Lives On

Your Data Lives On

- Stripe Size and Alignment
 - 512B or 4096B sectors?
 - Types/sizes of files being written
 - Make sure chunks are aligned to sector boundaries!

Performance!

Cool Features of MD

- Growing RAID devices
 - RAID 1 -> 5 -> 6
 - Replace drives with larger drives
 - Add additional drives

Conclusion

- Software RAID is great for SOHO uses where you can't justify a BBU RAID card
- Performance is typically good
- Hardware RAID can have better, more reliable storage due to BBU write cache

Doug Warner

doug@warner.fm