

INTERLAB

Tel: +331 3971 7337 – Fax: +331 3928 9008

ISSUE: xfs RAID OPTIMIZATION

Concerned platforms:

OS : Linux Fedora Core2 Kernel 2.6.x

Vernouillet, 17.10.05

Issue:

Getting the BEST performances on RAID

Tools:

- bech released by the conceptor of the application.
- Raid chassis with manual telnet access
- XFS filesystem

Action

Determin what kind of data is to be transported. HD, PAL or 2k video DPX or wrapped files...

Tweak point:

Hardware controller setup
Filesystem
Raidtool

Hardware controller setup

Tip 1, do not create a RAID5 volume till all parameters are clear. It takes ages, create a RAID0 and perform tests on it. Once the best stripe size is defined. Make your RAID5 volume with the right parameters.

```
+-----+
| Raid Level           : RAID 0
| Online SCSI Drives  : 7
| Maximum Drive Capacity : 285907 MB
| Disk Reserved Space  : 256 MB
| Logical Drive Assignment: Primary Controller
| Write Policy         : Default(Write-Back)
| Stripe Size         : 256K Bytes
+-----+

< Main Menu >
Quick installation
view
view Maximum Queued I/O Count - 512
```

INTERLAB

Tel: +331 3971 7337 - Fax: +331 3928 9008

```
view LUNs per Host SCSI ID - 8
view Max Number of Concurrent Host-LUN Connection - Def(4)
view Number of Tags Reserved for each Host-LUN Connection - 64
view Peripheral Device Type Parameters
v Host Cylinder/Head/Sector Mapping Configuration
s C Fibre Connection Option - Loop only
v C
v Host-side Parameters
  Drive-side Parameters
  Disk Array Parameters
  Controller Parameters
```

```
< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit Drives
view and edit channels
view and edit Configuration parameters
v ces
s Communication Parameters
v Caching Parameters
v H
  D Write-Back Cache Enabled
  D Optimization for Sequential I/O
  C Periodic Cache Flush Time - Disabled (Default)
```

BAT:NONE

```
< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit Drives
v
v CPU Type PPC750FX
v Total Cache Size 256MB DDR(ECC)
s Firmware Version 3.42A.05
v Bootrecord Version 1.22D
v FW Upgradability Rev. C
  Serial Number 6780135
  Battery Backup Unit Absent
  Base Board Rev. ID 0
  Base Board ID 81
  ID of NVRAM Defaults A16F-G2221 V6.06
  Controller Position Slot A
```

Generally, for big files bigger are stripes better it is. Other parameters are involved only testing will give you the magic setup.

INTERLAB

Tel: +331 3971 7337 – Fax: +331 3928 9008

Raidtool

Parameter & config file in /etc/raidtab

```
raiddev /dev/md0
raid-level          0
nr-raid-disks      2
persistent-superblock 1
chunk-size       512

device             /dev/sda
raid-disk          0
device             /dev/sdb
raid-disk          1
```

Explore chunk-size value 64 – 4096. do not forget to relaunch the raid with those values for this you must:

UNMOUNT the RaidVolume

```
[root@panoramix home]# umount /dev/md0
```

Make RAID

```
[root@panoramix home]# raidstop /dev/md0
[root@panoramix home]# mkraid /dev/md0
```

Test...

XFS tweaking:

Create an XFS file system

```
[root@panoramix home]# umount /dev/md0
[root@panoramix home]# mkfs.xfs -i size=2048 /dev/md0
```

for large filesystems, changing the inode size might help performance

This setup seems to be good for 2k DPX files. (Nocoda)

```
[root@panoramix home]# mkfs.xfs -d sectsize=4096 -l sectsize=4096
/dev/md0
```

This setup seems to be good for hd Panoramix.

INTERLAB

Tel: +331 3971 7337 – Fax: +331 3928 9008

NB:

Also, if you know number of disks and stripe, perhaps :

```
mkfs.xfs -d su=256k,sw=2 -i size=256 /dev/md0
```

Ceci est a essayer

Pls. send us your feel back.
Rev.1 – David 17.10.05