

Facultat d'Informàtica de Barcelona
Univ. Politècnica de Catalunya

Administració de Sistemes Operatius

Maintenance of the file system



Topics

- 1. Introduction to OS administration
- 2. Installation of the OS
- 3. Users management
- 4. Applications management
- 5. System monitoring
- 6. Maintenance of the file system
- 7. Local services
- 8. Network services
- 9. Protection and security

Objectives

- Knowledge
 - File systems
 - Backup tools
 - Backup support
- Abilities
 - Increase the size of a file system
 - Verify a file system
 - Make and restore backups

Filesystems (I)

- FAT (FAT16)
 - DOS filesystem
 - Small disks (< 4GB)
 - Filenames: 8+3
- FAT32 (VFAT)
 - Big disks
 - Long filenames
 - Direct links added
 - No owner nor permission control
- HPFS (NTFS)
 - Adds links and protections

Filesystems (II)

- ext2
 - UNIX filesystem
 - Soft/hard links
 - Access protection
 - Long filenames
- ext3
 - Adds journaling (eases error recovery)
- reiserfs
 - Organizes files and directories as a database
 - Supports journaling
 - Particularly effective with small files
 - Doesn't suffer from internal fragmentation

Filesystems (III)

- xfs
 - journaling
 - dynamic i-node management
 - ACLs
 - Huge maximum size
 - FS activity log
- jfs
 - journaling
 - dynamic i-node management
 - ACLs i MAC (Mandatory Access Control)
 - Huge maximum size

Journal filesystems

- *Journal*: log of all disk operations
 - Eases the FS recovery in case of a crash or error
 - Disk primitives are slightly slower
- The *journal* does not go through the *buffer cache*
 - The *journal* can be allocated in a different partition/disk
- Ext3, reiserfs, JFS, XFS all have a *journal*

Disk verification (I)

- Sources of problems
 - Hardware errors
 - Power outages
 - Operating system errors
 - Administration errors
 - Incorrect shutdown of the machine
- Never verify a filesystem while mounted/without syncing
 - High risk of corruption
 - The verification process access the device directly through the driver instead of going through the FS

Disk verification (II)

- Logical verification
 - Filesystem metadata
 - Directory structure
 - Lost data recovery
 - lost+found directory
- Physical verification
 - Disk blocks with I/O errors
 - **badblocks** command

Disk fragmentation

- Faster access to consecutive blocks of the disk
- Faster access to neighbour files in the disk
- Faster access depends on the disk area
 - Middle vs end

Extending the filesystem

- Install and configure a new disk
 - Partition it
 - Or recycle partitions in an old disk
- Decide the mount points
- Create the filesystem
- Transfer data to the new partition
- Mount the partition
 - Modify `/etc/fstab`
- We may need to reorganize the directories
 - `/home` -> `/homeA` + `/homeB`
 - `/home` -> `/home/alumnes` + `/home/professors`

Activity

- Plan and define possible extensions to the following directories
 - /home
 - /usr/local
 - /var

Other maintenance tasks

- Monitoring
 - Free space (df)
 - Most filesystem reserve some space for root use (5%)
 - Used space (du)
- Sync
 - Write to disk all *buffers* that have been modified



LVM

- Logical Volume Manager
 - High level abstraction of the storage space
 - Groups multiple physical partitions
 - new devices can be added to the volumes
 - Allows to define logical partitions
 - they can have logical names
 - they can be redistributed over the physical partitions
 - resize
 - move
 -


LVM

/dev/hda2 /dev/hda3 /dev/hdb1

Physical Volumes
(PV)



Logical Volumes
(LV)



LVM

- /etc/fstab:


/boot	/dev/hda1	...
swap	/dev/vg00/swap	...
/	/dev/vg00/root	...
/home	/dev/vg00/home	...
/usr	/dev/vg00/usr	...

Backup copies

- What data to copy?
 - User data (home, mails, ...)
 - Program data (Dbs, CVS, web, ...)
 - System configuration
 - Binaries?
- Copies frequency
 - Volatility of the data
 - Data significance
- Backup types
 - *Total Backup* (all)
 - *Incremental backup* (only changes)
 - *Reverse incremental backup* (only changes)


Total backup

- Everything is always copied
 - Fast to restore
 - Big size




Incremental Backup

- Only files that have changed are copied
 - From a certain date
- Slower to restore
- Small size
- First one is like a total backup
- A long incremental chain complicates the restore process



Reverse incremental backup

- Everything is copied, but in the previous backup only remains what has changed
 - Fast to restore
 - Small size
 - Only for random access devices (disks)



Backup copies (II)

- Physical support
 - Floppy, disk , CD, tape, network ..
 - We need to consider:
 - Availability
 - Usability
 - Speed
 - Cost/capacity
 - Liability
- Where do we keep the copies?
 - Accident protection
 - Fireproof boxes
 - Keep some out of the buildings
 - Theft protection
- Copy verification
- Data compression
 - Space vs security

Activity

- Define a backup policy (what to save, backup type, frequency, media, compression, ...) for:
 - A multiuser server for a company with:
 - 500 GB disk and 80 users
 - E-mail
 - 50MB per user
 - Web pages
 - 20 MB per user
 - 100 MB corporative web
 - Code repository
 - 10 GB distributed across 20 projects
 - Only 5 active projects

Other considerations

- Multiple servers
 - Use backup servers
 - more economics
 - easier to administrate
 - Tools: tar+rsync/ssh, amanda, bacula

Personal work

- Scheduling of tasks
 - Programming languages: sh, perl
 - Commands: find, grep...