




Directory Structure

- A collection of nodes containing information about all files



Both the directory structure and the files reside on disk




Disk Structure



- Disk can be subdivided into **partitions**
- Disks or partitions can be **RAID** protected against failure
- Disk or partition can be used **raw** – without a file system, or **formatted** with a file system
- Partitions also known as minidisks, slices
- Entity containing file system known as a **volume**
- Each volume containing file system also tracks that file system's info in **device directory** or **volume table of contents**
- As well as **general-purpose file systems** there are many **special-purpose file systems**, frequently all within the same operating system or computer



A Typical File-system Organization





Types of File Systems



- We mostly talk of general-purpose file systems
- But systems frequently have many file systems, some general- and some special- purpose
- Consider Solaris has
 - tmpfs – memory-based volatile FS for fast, temporary I/O
 - objfs – interface into kernel memory to get kernel symbols for debugging
 - cdfs – contract file system for managing daemons
 - lofs – loopback file system allows one FS to be accessed in place of another
 - procfs – kernel interface to process structures
 - ufs, zfs – general purpose file systems



Operations Performed on Directory



- Search for a file
- Create a file
- Delete a file
- List a directory
- Rename a file
- Traverse the file system



Directory Organization




The directory is organized logically to obtain

- Efficiency – locating a file quickly
- Naming – convenient to users
 - Two users can have same name for different files
 - The same file can have several different names
- Grouping – logical grouping of files by properties, (e.g., all Java programs, all games, ...)



Single-Level Directory

- A single directory for all users

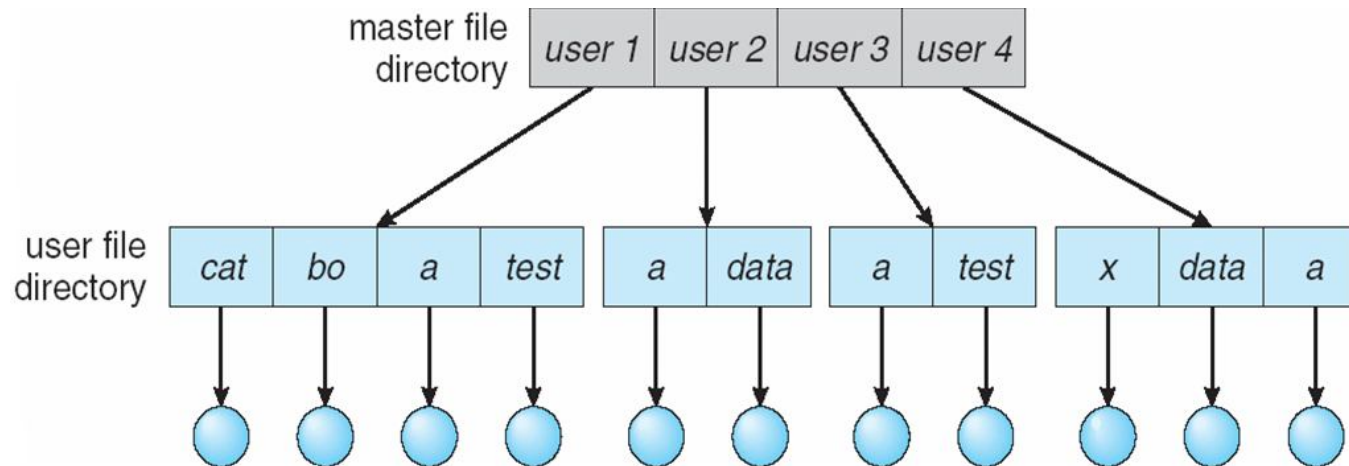


- Naming problem
- Grouping problem



Two-Level Directory

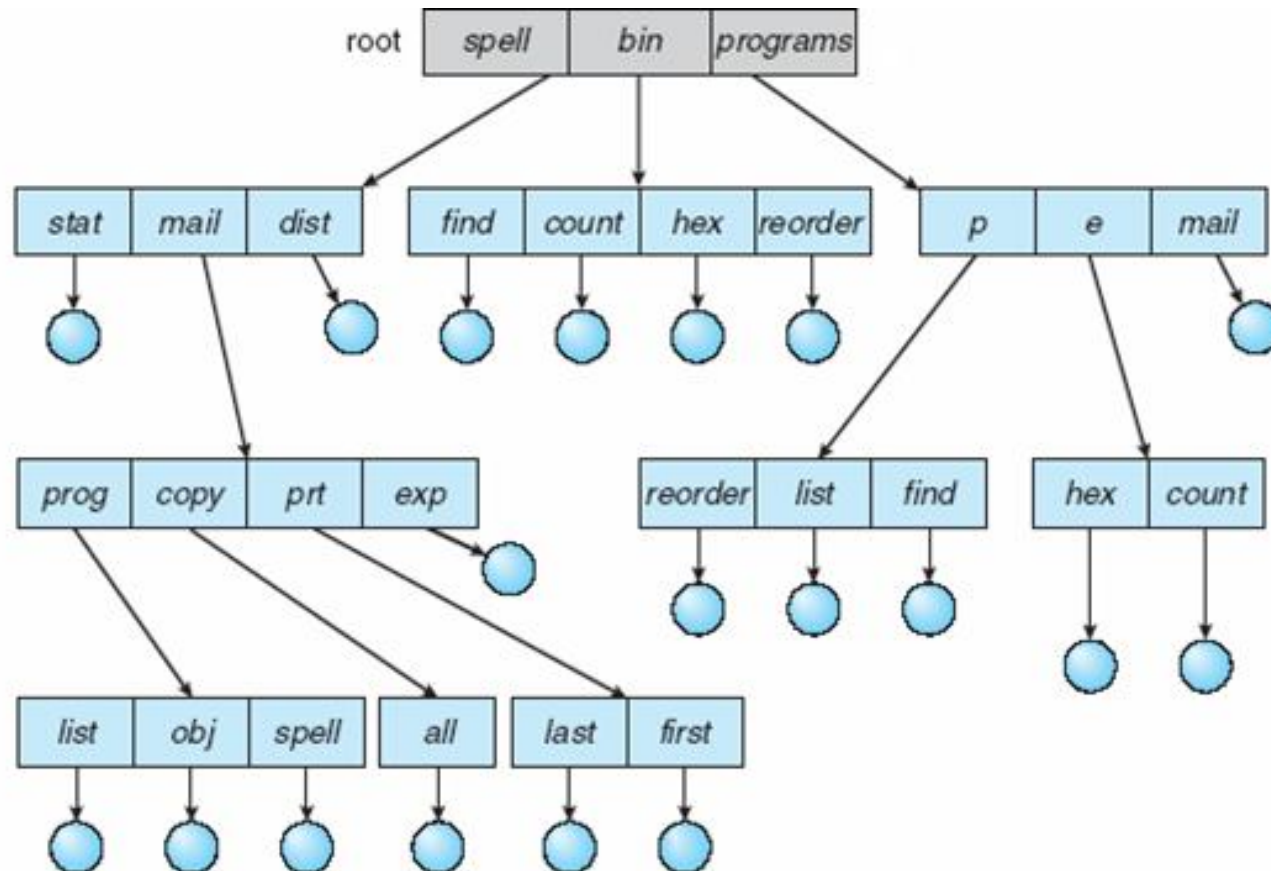
- Separate directory for each user



- Path name
- Can have the same file name for different user
- Efficient searching
- No grouping capability



Tree-Structured Directories





Tree-Structured Directories (Cont.)



- Efficient searching
- Grouping Capability
- Current directory (working directory)
 - `cd /spell/mail/prog`
 - `type list`



Tree-Structured Directories (Cont)



- **Absolute** or **relative** path name
- Creating a new file is done in current directory
- Delete a file


rm <file-name>

- Creating a new subdirectory is done in current directory

mkdir <dir-name>

Example: if in current directory **/mail**

mkdir count



Deleting "mail" ⇒ deleting the entire subtree rooted by "mail"



Acyclic-Graph Directories (Cont.)



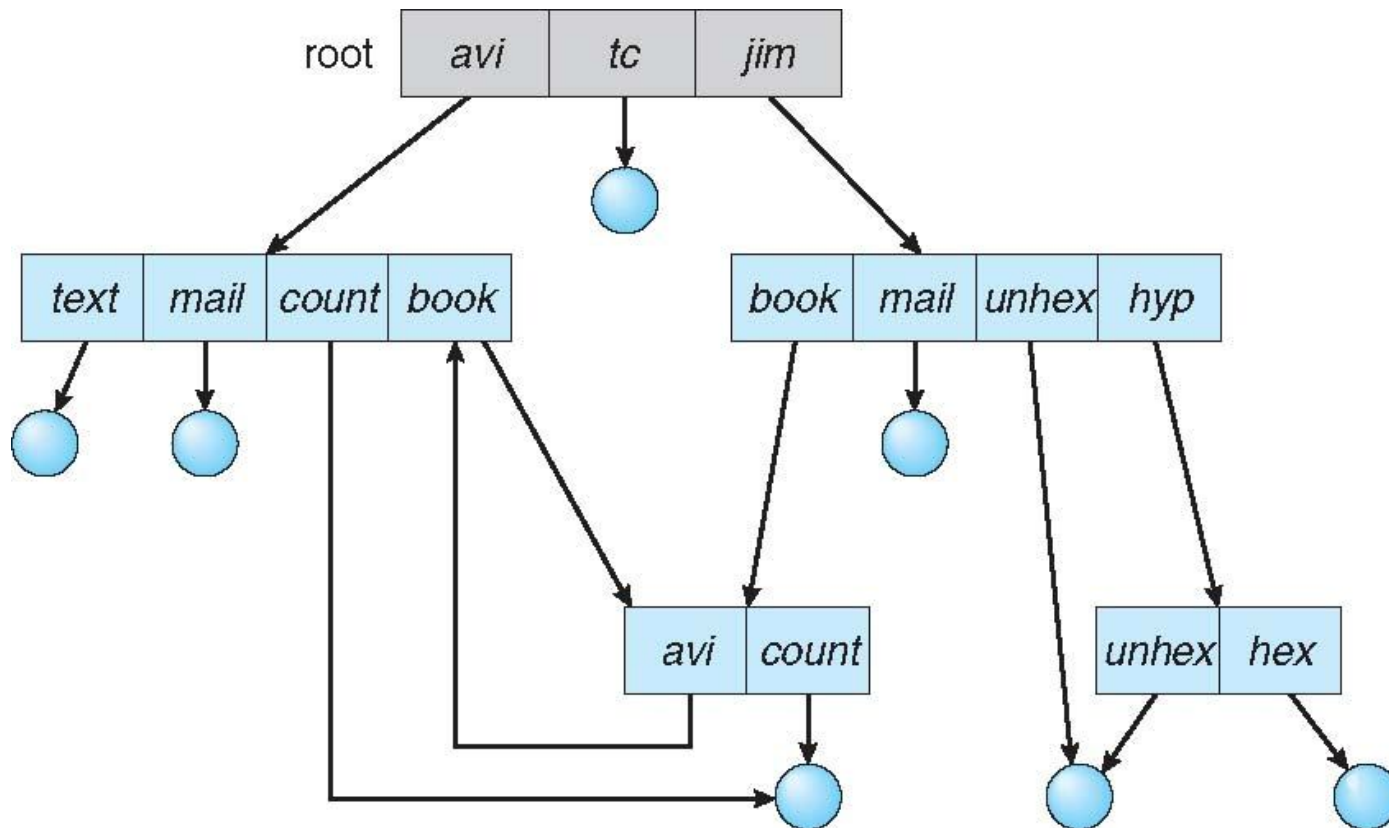
- Two different names (aliasing)
- If *dict* deletes *list* \Rightarrow dangling pointer

Solutions:

- Backpointers, so we can delete all pointers
Variable size records a problem
- Backpointers using a daisy chain organization
- Entry-hold-count solution
- New directory entry type
 - **Link** – another name (pointer) to an existing file
 - **Resolve the link** – follow pointer to locate the file



General Graph Directory





General Graph Directory (Cont.)




- How do we guarantee no cycles?
 - Allow only links to file not subdirectories
 - **Garbage collection**
 - Every time a new link is added use a cycle detection algorithm to determine whether it is OK




File System Mounting

- A file system must be **mounted** before it can be accessed
- A unmounted file system (i.e., Fig. 11-11(b)) is mounted at a **mount point**





Mount Point





File Sharing



- Sharing of files on multi-user systems is desirable
- Sharing may be done through a **protection** scheme
- On distributed systems, files may be shared across a network
- Network File System (NFS) is a common distributed file-sharing method
- If multi-user system
 - **User IDs** identify users, allowing permissions and protections to be per-user
 - Group IDs** allow users to be in groups, permitting group access rights
 - Owner of a file / directory
 - Group of a file / directory



File Sharing – Remote File Systems



- Uses networking to allow file system access between systems
 - Manually via programs like FTP
 - Automatically, seamlessly using **distributed file systems**
 - Semi automatically via the **world wide web**
- **Client-server** model allows clients to mount remote file systems from servers
 - Server can serve multiple clients
 - Client and user-on-client identification is insecure or complicated
 - **NFS** is standard UNIX client-server file sharing protocol
 - **CIFS** is standard Windows protocol
 - Standard operating system file calls are translated into remote calls
- Distributed Information Systems (**distributed naming services**) such as LDAP, DNS, NIS, Active Directory implement unified access to information needed for remote computing



File Sharing – Failure Modes



- All file systems have failure modes
 - For example corruption of directory structures or other non-user data, called **metadata**
- Remote file systems add new failure modes, due to network failure, server failure
- Recovery from failure can involve **state information** about status of each remote request
- **Stateless** protocols such as NFS v3 include all information in each request, allowing easy recovery but less security



File Sharing – Consistency Semantics



- Specify how multiple users are to access a shared file simultaneously
 - Similar to Ch 5 process synchronization algorithms
 - Tend to be less complex due to disk I/O and network latency (for remote file systems)
 - Andrew File System (AFS) implemented complex remote file sharing semantics
 - Unix file system (UFS) implements:
 - Writes to an open file visible immediately to other users of the same open file
 - Sharing file pointer to allow multiple users to read and write concurrently
 - AFS has session semantics
 - Writes only visible to sessions starting after the file is closed



Protection



- File owner/creator should be able to control:
 - what can be done
 - by whom
- Types of access
 - **Read**
 - **Write**
 - **Execute**
 - **Append**
 - **Delete**
 - **List**




Access Lists and Groups



- Mode of access: read, write, execute
- Three classes of users on Unix / Linux

a) owner access	7	⇒	RWX 111
b) group access	6	⇒	RWX 110
c) public access	1	⇒	RWX 001

- Ask manager to create a group (unique name), say G, and add some users to the group.
- For a particular file (say *game*) or subdirectory, define an appropriate access.

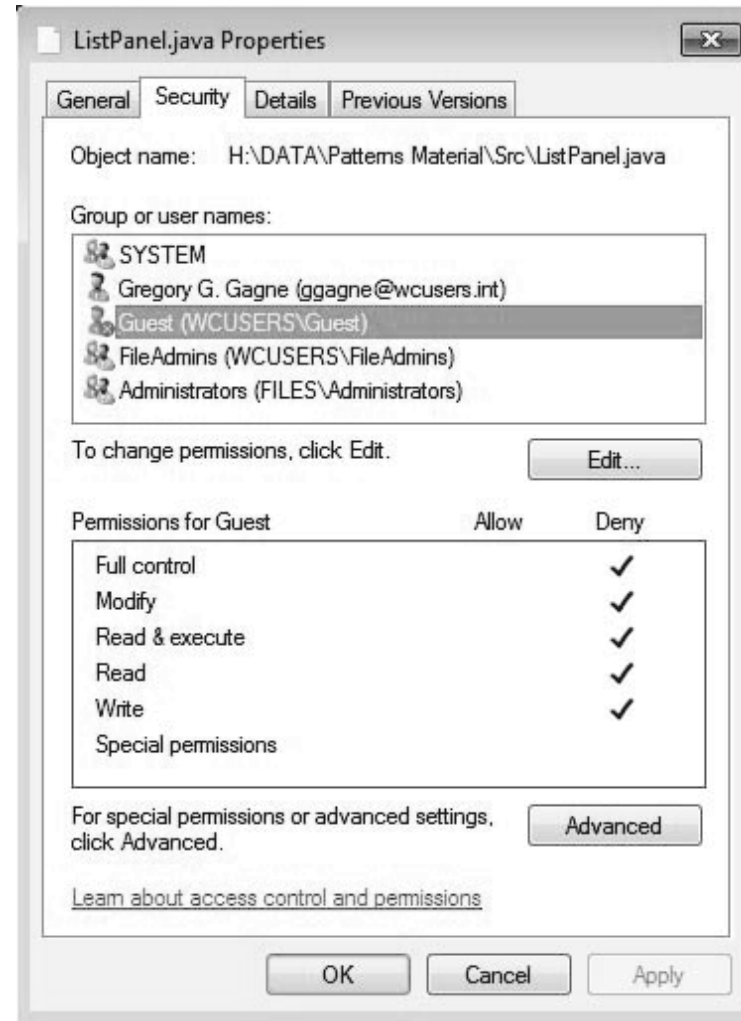


Attach a group to a file

chgrp **G** **game**



Windows 7 Access-Control List Management





A Sample UNIX Directory Listing



```
-rw-rw-r--  1 pbg  staff    31200  Sep 3 08:30  intro.ps
drwx-----  5 pbg  staff     512  Jul 8 09:33  private/
drwxrwxr-x  2 pbg  staff     512  Jul 8 09:35  doc/
drwxrwx---  2 pbg  student   512  Aug 3 14:13  student-proj/
-rw-r--r--  1 pbg  staff    9423  Feb 24 2003  program.c
-rwxr-xr-x  1 pbg  staff   20471  Feb 24 2003  program
drwx--x--x  4 pbg  faculty   512  Jul 31 10:31  lib/
drwx-----  3 pbg  staff    1024  Aug 29 06:52  mail/
drwxrwxrwx  3 pbg  staff     512  Jul 8 09:35  test/
```