

CSE 451: Operating Systems Spring 2006

Module 12.5 Secondary Storage Overview

John Zahorjan
zahorjan@cs.washington.edu
Allen Center 534

Interface Layers

5/17/2006 © 2006 Gribble, Lazowska, Levy, Zahorjan 2

Exported Abstractions

5/17/2006 © 2006 Gribble, Lazowska, Levy, Zahorjan 3

Primary Roles of the OS (file system)

1. Hide hardware specific interface
2. Allocate disk blocks
3. Check permissions
4. Understand directory file structure
5. Maintain metadata
6. Performance
7. Flexibility

*Why does the OS define directories?
Why not leave that to the library/application layer?
(Why would you want to leave it to the app/library?)*

5/17/2006 © 2006 Gribble, Lazowska, Levy, Zahorjan 4

What Is A File?

- At this level, what we think of as a file is an "inode"
- An inode keeps track of where the data blocks are
- It also keeps track of metadata (e.g., permissions, last modification time, etc.)

5/17/2006 © 2006 Gribble, Lazowska, Levy, Zahorjan 5

What Is A Directory?

Ok, so what is a "hard link"?

17	..
2	..
21812	passwd
45700	hosts

5/17/2006 © 2006 Gribble, Lazowska, Levy, Zahorjan 6

Obtaining Performance: Caching

read write → OS → readblock writeblock → Disk

No...

Block in cache?

- Why cache?
- How does it affect applications?

5/17/2006 © 2006 Gribble, Lazowska, Levy, Zahorjan 7

Obtaining Performance: open

```
int fd = open("/root/xxx", O_RDONLY)
[ read(fd,...)
  write(fd,...) ]
```

Per-process open file table

0
1
2
3
...

file pointer

inode

root etc

5/17/2006 © 2006 Gribble, Lazowska, Levy, Zahorjan 8

Multiple File Systems

read write → [ext2, fat] → readblock writeblock → ext2 disk / fat disk

Each filesystem:

- Decides which blocks to allocate for what (inodes, directories, file data)
- Decides what the format of a directory is

Not really disks, but partitions. We'll get to that later...

5/17/2006 © 2006 Gribble, Lazowska, Levy, Zahorjan 9

Supporting Multiple File Systems: vfs

read write → [vfs] → [ext2, fat] → ext2 disk / fat disk

5/17/2006 © 2006 Gribble, Lazowska, Levy, Zahorjan 10

Extending Functionality

read write → [vfs] → [ext2, fat, /dev, /proc] → ext2 disk / fat disk

- /dev: direct access to hardware blocks
- /proc: "virtual files" (computed, not stored, results)
 - /proc is to the file system what cgi's are to a web server (sort of)

5/17/2006 © 2006 Gribble, Lazowska, Levy, Zahorjan 11

Are Directories Fundamental?

What is the logical role of the directories?

What practical role(s) do they have?

Are there alternatives?

5/17/2006 © 2006 Gribble, Lazowska, Levy, Zahorjan 12