

Administrivia

- Homework 5 due today.
- Homework 6 due Tuesday.
- Review sheet for exam on Web by Monday.

Slide 1

Journaling Filesystems — Overview

- Recall — o/s sometimes doesn't perform "write to disk" operations right away (caching).
- One result is likely improved performance. Another is potential filesystem inconsistency — operations such as "move a block from the free list to a file" are no longer atomic.
- Idea of journaling filesystem — do something so we *can* regard updates to filesystem as atomic.
- To say it another way — record changes-in-progress in log, when complete mark them "done".

Slide 3

Disk Fragmentation

- Idea — if blocks that make up a file are (mostly) contiguous, faster to read them all. If not, "disk fragmentation".
- How likely is disk fragmentation? Depends on filesystem, strategy for allocating space for files.
- "Defragmenter" utility can be run to correct it. Windows comes with one. Linux doesn't. The claim is that Unix and Linux filesystems typically don't become fragmented unless the disk is close to full.

Slide 2

Journaling Filesystems, Continued

- Can record "data", "metadata" (directory info, free list, etc.), or both.
- "Undo logging" versus "redo logging":
 - Undo logging: First copy old data to log, then write new data (possibly many blocks) to disk. If something goes wrong during update, "roll back" by copying old data from log.
 - Redo logging: First write new data to log (i.e., record changes we're going to make), then write new data to disk. If something goes wrong during update, complete the update using data in log.
- A key benefit — after a system crash, we should only have to look at the log for incomplete updates, rather than doing a full filesystem consistency check.

Slide 4

Journaling Filesystems Versus Log-Structured Filesystems

- Log-structured filesystem — *everything* is written to log, and only to log. Seems like an interesting idea, but tough to implement with good performance.
- Journaling filesystem — log contains only recent and pending updates.

Slide 5

Unix Filesystems — Concepts

- Single type of file — sequence of bytes. `lseek` allows random access.
- Single root directory. `mount` allows access to multiple physical devices.
- Links, hard or symbolic, to allow non-tree directory structure.
- Locks to control access to files/records.
- File descriptors for open files.
- "Pipes" — pseudofiles for connecting processes.

Slide 6

Unix Filesystems — Implementation Overview

- Superblock contains critical info — how many i-nodes, location of free list, how many blocks, etc.
- After that? In early implementations, all i-nodes followed by all data blocks. Later implementations (Berkeley FFS) use "cylinder groups", each containing superblock copy, i-nodes, and data blocks — for better performance, reliability.
- Directory entries fixed-size in early implementations, varying-size in later ones.
- I-nodes contain file attributes, link count, list of (some) blocks, pointers to indirect blocks.
- In memory — table of i-nodes for open files, table of file descriptors (containing, e.g., info about position within file).
- NFS allows access to other systems' disks.

Slide 7

Linux Filesystems — Implementation Overview

- Originally, MINIX filesystem only — similar to early Unix.
- Later, VFS (virtual filesystem) added as intermediate layer to support many kinds.
- ext2/ext3 filesystems (ext3 is ext2 with addition of "journal" file and support for journaling): Similar to FFS, but with a single blocksize, and "block groups" rather than "cylinder groups". Block group also includes bitmap for free space. Attempts to allocate all space for file within block group (may account for less fragmentation). Superblock has bit that says whether filesystem is "clean" (no `fsck` needed at boot time).
- `/proc` filesystem represents much system info.

Slide 8

Windows Filesystems — Concepts

- FAT filesystems as described in chapter 6. Designed for small disks and don't work very well for large ones. Also see description of how support for long filenames was added.
- NTFS filesystems — new for Windows NT. Some basic concepts:
 - Unicode filenames.
 - File can consist of multiple "streams" (not just one as in Unix) — generalization of Mac's data fork / resource fork idea.
 - Transparent compression and encryption.

Slide 9

Minute Essay

- I'll review for exam Tuesday. Anything you particularly want to hear about? Or send me mail later.

Slide 11

Windows Filesystems — NTFS Implementation

- MFT (master file table) — analogous to i-nodes. One or more entries per file/directory, plus some for system. Boot block points to start.
- For small files, data is right in MFT record. Otherwise contains list of contiguous sequences of bytes.

Slide 10