

File Systems

- Implement the file abstraction
 - In Unix and Windows
 - File == sequence of bytes
 - Directories == Nodes for accessing files

On-disk file formats

- Metadata blocks
 - Boot blocks / Partition tables
 - File system information (Superblocks / Cylinder group summaries)
- File control blocks (inodes in Unix)
- Directory blocks
- Data blocks
- Log records (in modern implementations)

Unix “Fast” File system

- Inodes
 - 256-byte records stored in cylinder groups
 - Inode #2 is the root of the file system
 - Contents include
 - File permissions
 - File access/modification/inode change times
 - Disk block access of first 12 records
 - Block address of “index” blocks for larger files
- Directories
 - Lists of <file name, inode number> pairs
- Cylinder groups are used to optimize file/directory access

IBM Journal File System

- Uses log disk to record “transactions”
- To create a new file
 - Write log record
 - Continue the applications
 - Make FFS disk updates at needed

MS/DOS FAT

- File Allocation Table
 - One entry per cluster
 - Links the clusters used by a file
- Directories
 - 32-byte entries
 - File name (8.3 DOS name)
 - File attribute byte
 - Directory/Hidden/System/Read-only
 - Last change time
 - File size
 - Starting cluster
 - About long file names
 - One entry contains the short alias
 - PROGRA~1
 - Additional entries (with weird attribute bytes) contain long name
 - Program Files
- Root directory
 - Stored after two copies of the FAT (pre FAT32)

Serialization and Consistency

One file system operation may require several I/O operations
Must sequential I/O be performed serially?
If not, how is consistency achieved

More modern file system

AdvFS (Advanced File System) in Compaq Tru64 Unix
Microsoft NTFS
Linux [ReiserFS](#)

Files are allocated (if possible) in large extents

AdvFS: extent map
NTFS: run-information

File control information is stored in flexible structures

AdvFS: BMT (Bitfile Metadata Table)
NTFS: MFT (Master File Table)

Directory structures support larger file names and larger directories

AdvFS: "Old" FFS format with B+ trees
NTFS: Directories contain long file names and uses tree-structure for fast access

Logging with redo/undo records