

# Operating Systems

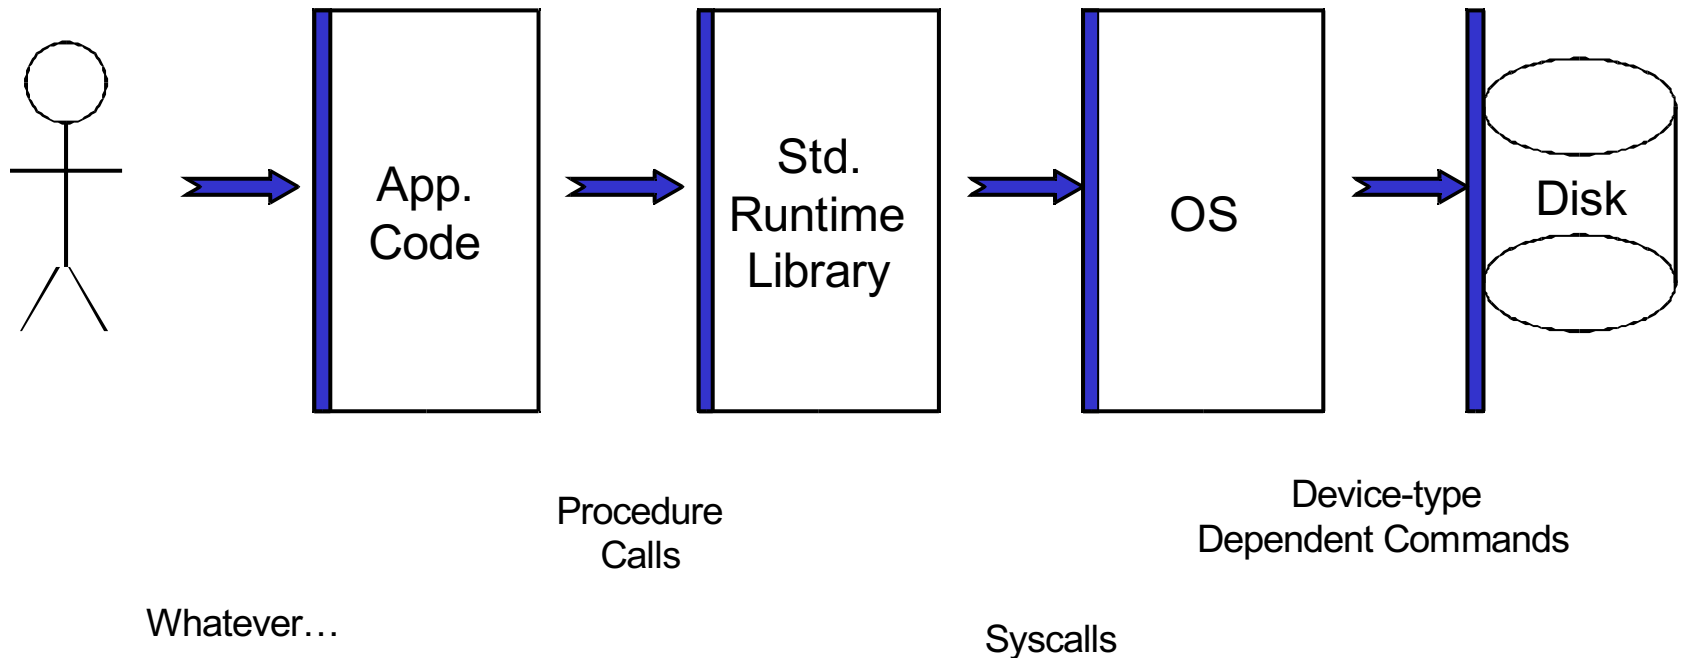
## File Systems

Lecture 13  
Michael O'Boyle

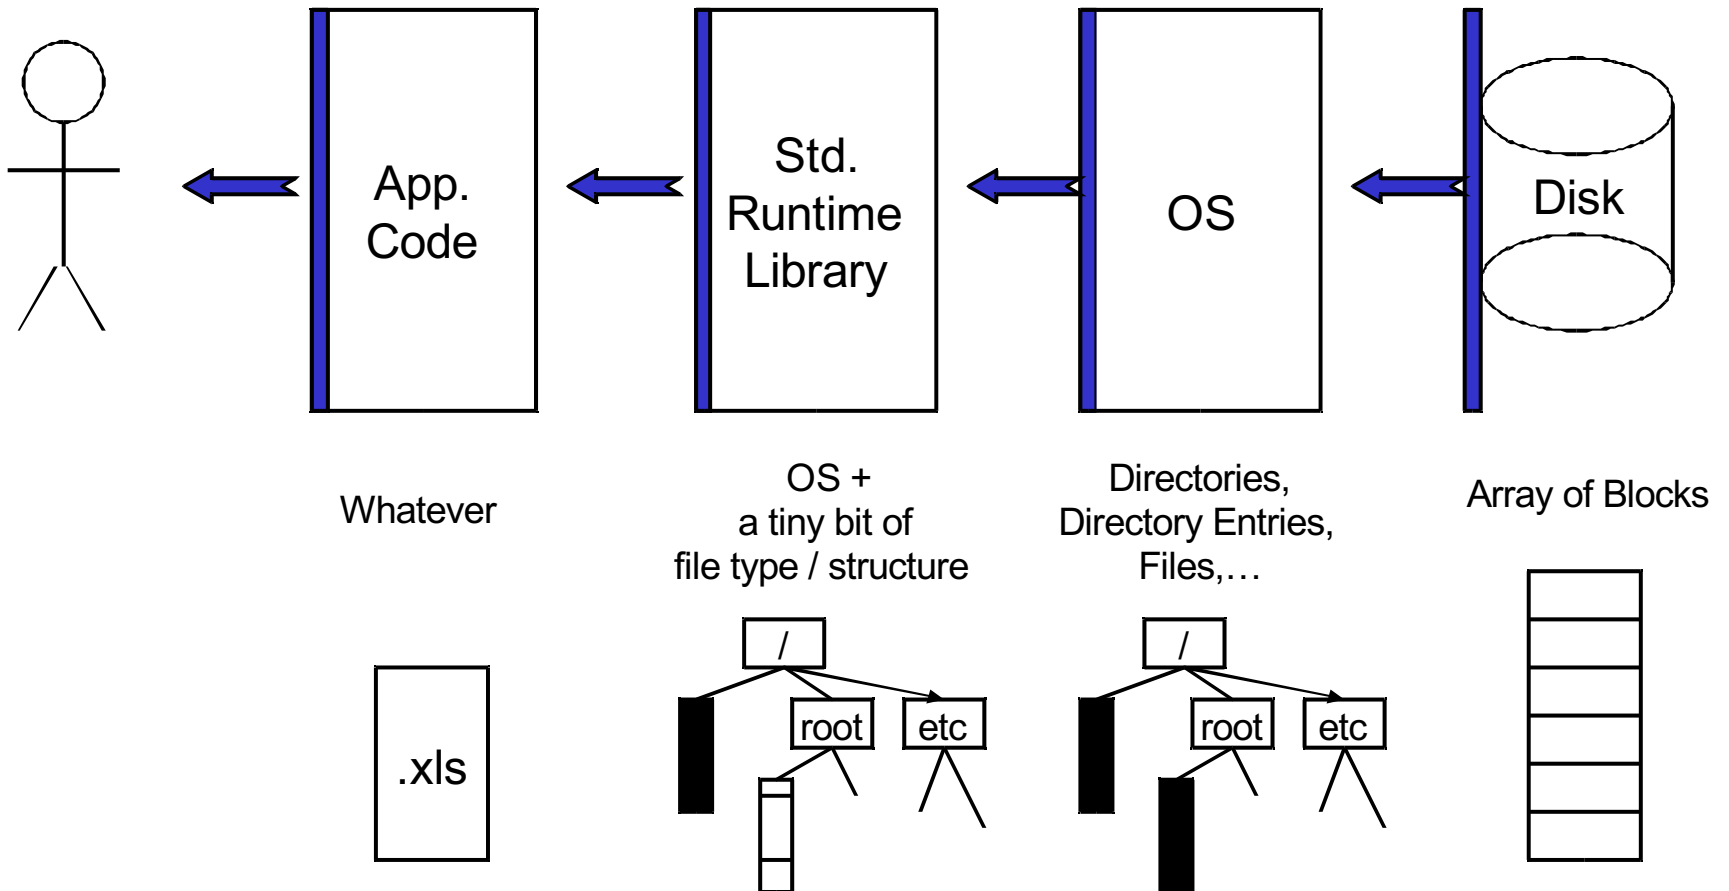
# Overview

- Roles
- Files
- Directories
- File Protection
- Unix
  - inodes

# Interface Layers




# Exported Abstractions



# 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?)*



# File systems

- The concept of a file system is simple
- The implementation of the abstraction for secondary storage
  - abstraction = files
- Logical organization of files into directories
  - the directory hierarchy
- Sharing of data between
  - processes, people and machines
- Access control, consistency, ...

# Files

- A collection of data with some properties
  - contents, size, owner, last read/write time, protection ...
- Files types
  - understood by file system
    - device, directory, symbolic link
  - understood by other parts of OS or by runtime libraries
    - executable, dll, source code, object code, text file, ...
- Type can be encoded in the file's name or contents
  - windows encodes type in name
    - .com, .exe, .bat, .dll, .jpg, .mov, .mp3, ...
- Old Mac OS stored the name of the creating program along with the file
- Unix - initial characters (e.g., #)

# Basic operations

## Unix

- create(name)
- open(name, mode)
- read(fd, buf, len)
- write(fd, buf, len)
- sync(fd)
- seek(fd, pos)
- close(fd)
- unlink(name)
- rename(old, new)

## Windows

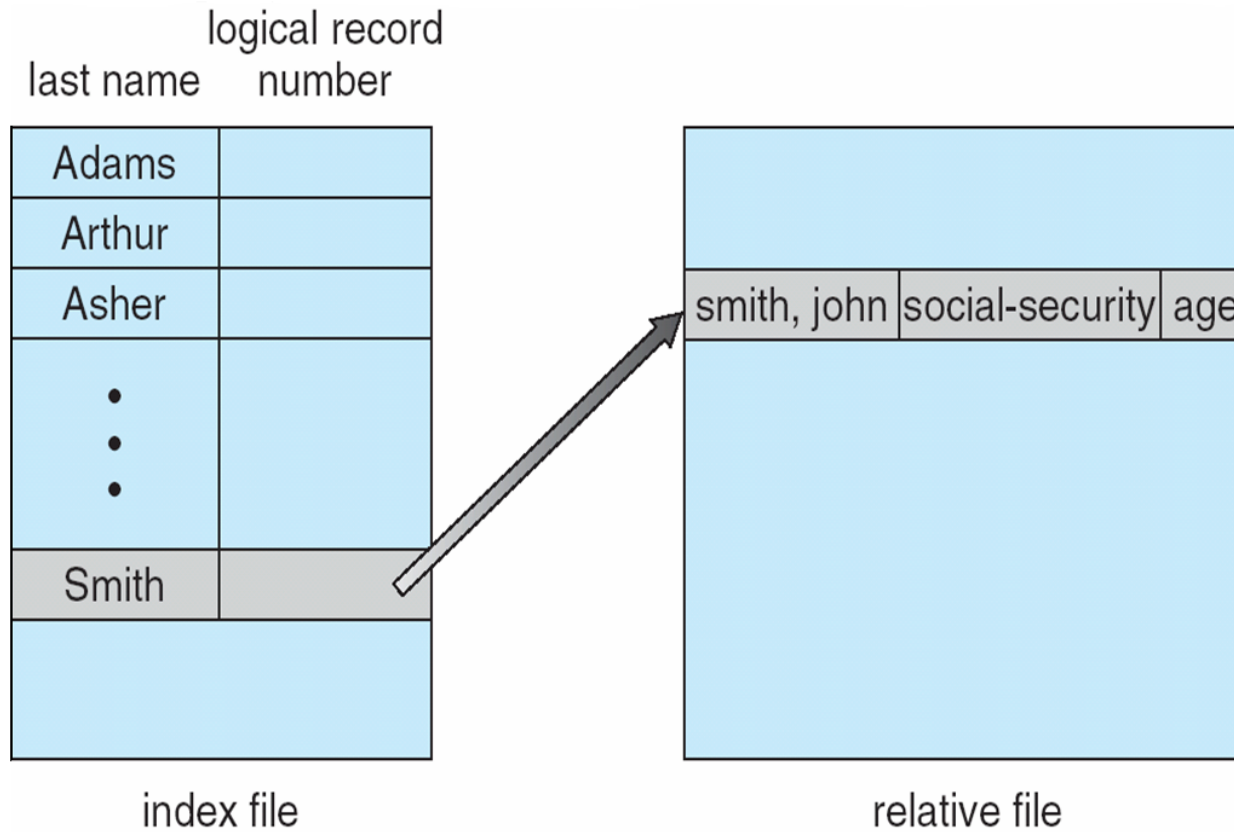
- CreateFile(name, CREATE)
- CreateFile(name, OPEN)
- ReadFile(handle, ...)
- WriteFile(handle, ...)
- FlushFileBuffers(handle, ...)
- SetFilePointer(handle, ...)
- CloseHandle(handle, ...)
- DeleteFile(name)
- CopyFile(name)
- MoveFile(name)



# File access methods

- File systems provide different **access methods**
  - sequential access
    - read bytes one at a time, in order
  - direct access
    - random access given a block/byte #
  - record access
    - file is array of fixed- or variable-sized records
  - indexed access
    - One file contains an index to a record in another file
    - apps can find a file based on value in that record (similar to DB)
- Why do we care about distinguishing sequential from direct access?
  - what might the FS do differently in these cases?

# Example of Index and Relative Files



# Directories

- Directories provide:
  - a way for users to organize their files
  - a convenient file name space for both users and FS's
- Most file systems support multi-level directories
  - naming hierarchies (`/`, `/usr`, `/usr/local`, `/usr/local/bin`, ...)
- Most file systems support the notion of current directory
- Absolute names: fully-qualified starting from root of FS

```
bash$ cd /usr/local
```
- Relative names: specified with respect to current directory

```
bash$ cd /usr/local    (absolute)
bash$ cd bin           (relative, equivalent to cd /usr/local/bin)
```

# Directory internals

- A directory is typically just a file that happens to contain special metadata
- Directory
  - list of (name of file, file attributes)
- Attributes include such things as:
  - size, protection, location on disk, creation time, access time, ...
- The directory list is usually unordered (effectively random)
  - when you type “ls”, the “ls” command sorts the results for you

# Path name translation

- You want to open “/one/two/three”

```
fd = open("/one/two/three", O_RDWR);
```

- Inside the file system

- open directory “/” (well known, can always find)
- search the directory for “one”, get location of “one”
- open directory “one”, search for “two”, get location of “two”
- open directory “two”, search for “three”, get loc. of “three”
- open file “three”
- (of course, permissions are checked at each step)

- FS spends much time walking down directory paths
  - this is why open is separate from read/write (session state)
  - OS will cache prefix lookups to enhance performance
    - /a/b, /a/bb, /a/bbb all share the “/a” prefix

# File protection

- FS implements a protection system
  - to control who can access a file (user)
  - to control how they can access it (e.g., read, write, or exec)
- Often generalised:
  - generalize files to **objects** (the “what”)
  - generalize users to **principals** (the “who”, user or program)
  - generalize read/write to **actions** (the “how”, or operations)
- Protection system dictates
  - whether a given action
  - performed by a given principal
  - on a given object should be allowed
- E.g., you can read or write your files, but others cannot
- E.g., your can read `/group/teaching/cs3` but you cannot write to it

# Model for representing protection

- Two different ways of thinking about it:
  - access control lists (ACLs)
    - for each object, keep list of principals and principals' allowed actions
  - capabilities
    - for each principal, keep list of objects and principal's allowed actions
- Both can be represented with the following matrix:

	/etc/passwd	/home/gribble	/home/guest
root	rw	rw	rw
gribble	r	rw	r
guest			r

objects

principals

ACL

capability

# ACLs vs. Capabilities

- Capabilities are easy to transfer
  - they are like keys: can hand them off
  - they make sharing easy
- ACLs are easier to manage
  - object-centric, easy to grant and revoke
    - to revoke capability, need to keep track of principals that have it
    - hard to do, given that principals can hand off capabilities
- ACLs grow large when object is heavily shared
  - can simplify by using “groups”
    - put users in groups, put groups in ACLs
- Additional benefit
  - change group membership, affects ALL objects that have this group in its ACL



# The original Unix file system

- Dennis Ritchie and Ken Thompson, Bell Labs, 1969
- “UNIX rose from the ashes of a multi-organizational effort in the early 1960s to develop a dependable timesharing operating system” – Multics
- Designed for a “workgroup” sharing a single system
- Did its job well
  - Although it has been stretched in many directions




# (Old) Unix disks are divided into five parts ...

- **Boot block**
  - can boot the system by loading from this block
- **Superblock**
  - specifies boundaries of next 3 areas, and contains head of freelists of inodes and file blocks
- **inode area**
  - contains descriptors (inodes) for each file on the disk; all inodes are the same size; head of freelist is in the superblock
- **File contents area**
  - fixed-size blocks; head of freelist is in the superblock
- **Swap area**
  - holds processes that have been swapped out of memory

# Basic file system structures

- Every file and directory is represented by an inode
  - Stands for “index node”
- Contains two kinds of information:
  - 1) Metadata describing the file's owner, access rights, etc.
  - 2) Location of the file's blocks on disk



# inode format

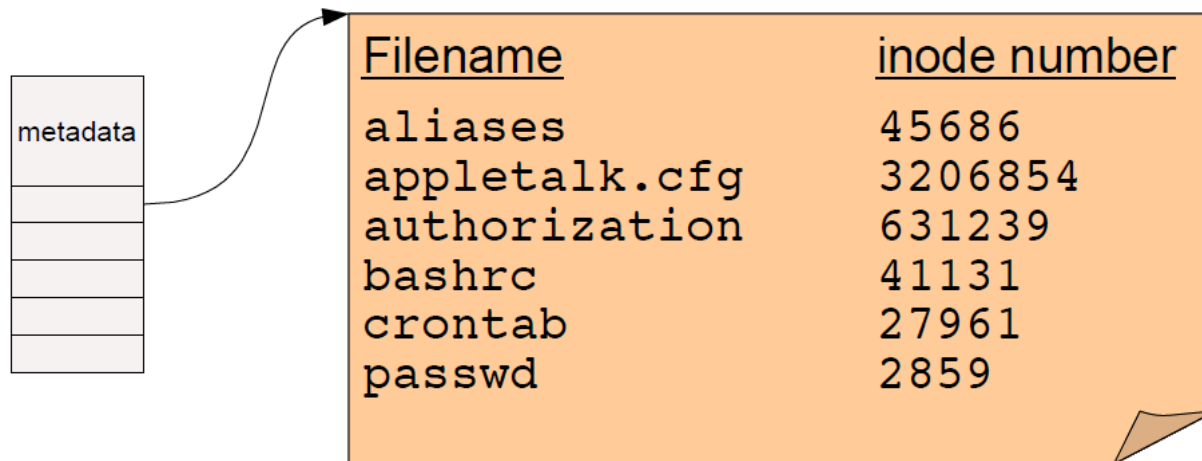
- User number
- Group number
- Protection bits
- Times (file last read, file last written, inode last written)
- File code: specifies if the inode represents a directory, an ordinary user file, or a “special file” (typically an I/O device)
- Size: length of file in bytes
- Block list: locates contents of file (in the file contents area)
- Link count: number of directories referencing this inode

# The flat (inode) file system

- Each file is known by a number, which is the number of the inode
  - 1, 2, 3, etc.!
- Files are created empty, and grow when extended through writes

# The tree (directory, hierarchical) file system

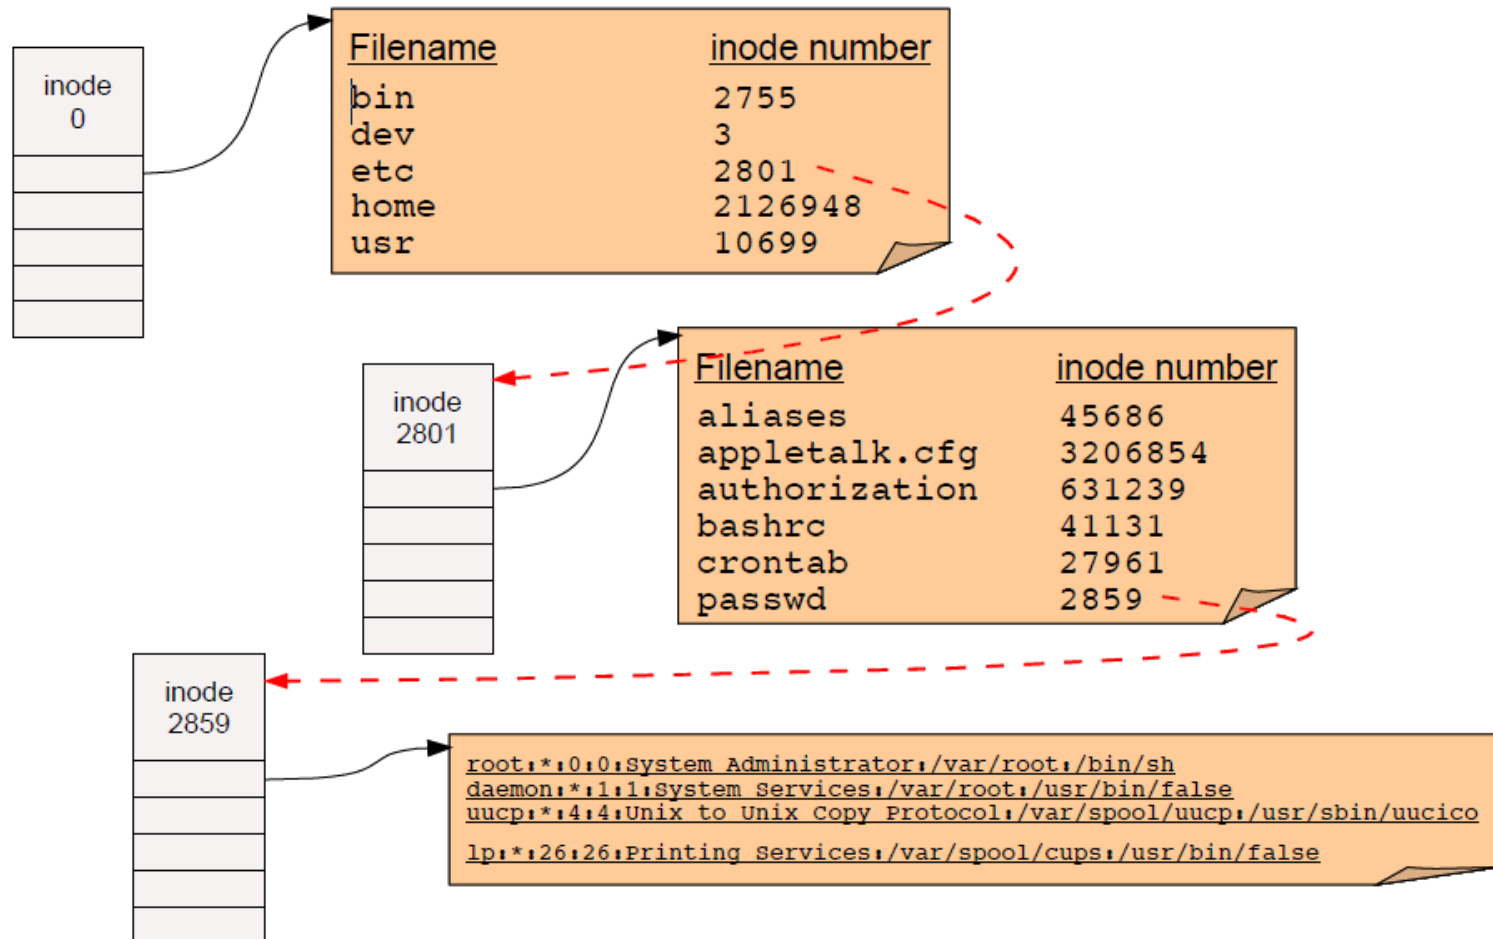
- A directory is a flat file of fixed-size entries
- Each entry consists of an inode number and a file name



- These are the contents of the directory “file data” itself – NOT the directory's inode
- Filenames (in UNIX) are not stored in the inode at all
- Special inodes:
  - inode 2 is the root directory
  - Inode 1 – a hidden file containing all bad blocks


# Pathname resolution

- To look up a pathname “/etc/passwd”, start at root directory and walk down chain of inodes...



# Locating inodes on disk

- Directories give the **inode number** of a file.
  - How do we find the inode itself on disk?
- Basic idea: Top part of filesystem contains *all* of the inodes



- inode number is just the “index” of the inode
- Easy to compute the block address of a given inode:
  - $\text{block\_addr}(\text{inode\_num}) = \text{block\_offset\_of\_first\_inode} + (\text{inode\_num} * \text{inode\_size})$
  - This implies that a filesystem has a fixed number of potential inodes
    - This number is generally set when the filesystem is created
  - The superblock stores important metadata on filesystem layout, list of free blocks, etc.



# Directory issues

- Directories map filenames to inode numbers.
- We can create multiple pointers to the same inode in different directories
  - Or even the same directory with different filenames
- In UNIX this is called a “hard link” and can be done using “ln”

```
bash$ ls -i /home/foo
```

```
287663 /home/foo          (This is the inode number of “foo”)
```

```
bash$ ln /home/foo /tmp/foo
```

```
bash$ ls -i /home/foo /tmp/foo
```


```
287663 /home/foo
```

```
287663 /tmp/foo
```


- “/home/foo” and “/tmp/foo” now refer to the same file on disk
  - Not a copy! You will always see identical data no matter which filename you use to read or write the file.

# How should we organize blocks on a disk?

- Very simple policy: A file consists of linked blocks
  - inode points to the first block of the file
  - Each block points to the next block in the file (just a linked list on disk)




- Indexed files
  - inode contains a list of block numbers containing the file
  - Array is allocated when the file is created



# The “block list” portion of the inode (Unix Version 7)

- Points to blocks in the file contents area
- Must be able to represent very small and very large files.
- Each inode contains 13 block pointers
  - first 10 are “direct pointers” (pointers to 512B blocks of file data)
  - then, single, double, and triple indirect pointers



# File Size

- Only occupies  $13 \times 4\text{B}$  in the inode
- Can get to  $10 \times 512\text{B} = \text{a } 5120\text{B}$  file directly
  - (10 direct pointers, blocks in the file contents area are 512B)
- Can get to  $128 \times 512\text{B} = \text{an additional } 65\text{KB}$  with a single indirect reference
  - (the 11<sup>th</sup> pointer in the inode gets you to a 512B block in the file contents area that contains 128 4B pointers to blocks holding file data)
- Can get to  $128 \times 128 \times 512\text{B} = \text{an additional } 8\text{MB}$  with a double indirect reference
  - (the 12<sup>th</sup> pointer in the inode gets you to a 512B block in the file contents area that contains 128 4B pointers to 512B blocks in the file contents area that contain 128 4B pointers to 512B blocks holding file data)

# File Size

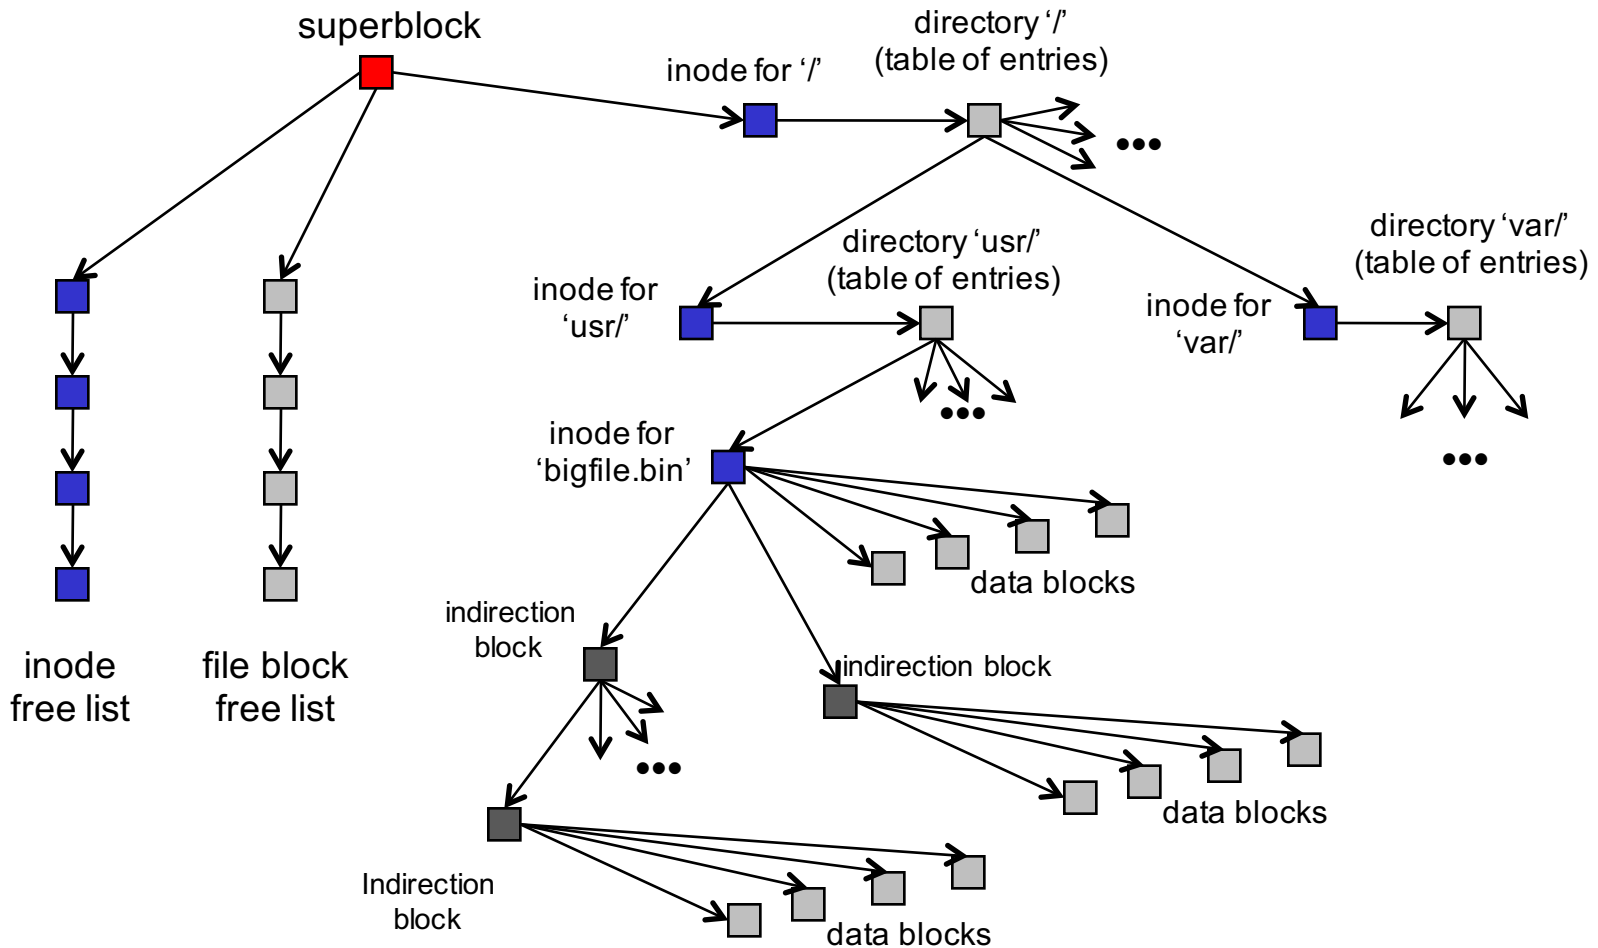
- Can get to  $128 \times 128 \times 128 \times 512\text{B} =$  an additional 1GB with a triple indirect reference
  - (the 13<sup>th</sup> pointer in the inode gets you to a 512B block in the file contents area that contains 128 4B pointers to 512B blocks in the file contents area that contain 128 4B pointers to 512B blocks in the file contents area that contain 128 4B pointers to 512B blocks holding file data)
- Maximum file size is a little over 1GB +

# Updated File Size

- A later version of Bell Labs Unix utilized 12 direct pointers rather than 10
- Berkeley Unix went to 1KB block sizes
  - What's the effect on the maximum file size?
    - $256 \times 256 \times 256 \times 1K = 17 \text{ GB} + \text{a smidge}$
- Subsequently went to 4KB blocks
  - $1K \times 1K \times 1K \times 4K = 4TB$

# Putting it all together

- The file system is just a huge data structure



# File system layout

- One important goal of a filesystem is to lay this data structure out on disk
  - keep in mind the physical characteristics of the disk
  - seeks are expensive
  - characteristics of the workload
    - locality across files within a directory,
    - sequential access to many files
- Old UNIX's layout is very inefficient
  - constantly seeking
    - between inode area and data block area
    - as traverse filesystem, or sequentially read files
- Newer file systems are smarter
- Newer storage devices (SSDs) change the constraints



# Summary

- Roles
- Files
- Directories
- File Protection
- Unix
  - Inodes
- Disks next