

## Section 9:

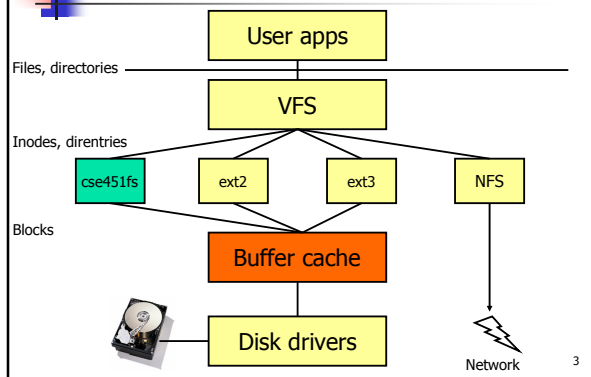
### Project 3 – The Buffer Cache Network File Systems

1

## Questions?

2

## Linux FS Layers (Revisit)



3

## Linux Buffer Manager

- Buffer cache caches disk blocks & buffers writes
  - When you write to a file, the data goes into buffer cache (for write-back buffer caches)
  - Sync is required to flush the data to disk
    - Update and bdflush processes flush data to disk (every 30s)
- Linux buffer cache code fundamentals:
  - Blocks are represented by *buffer\_heads*
    - Actual data is in *buffer\_head->b\_data*
  - For a given disk block, buffer manager could be:
    - Completely unaware of it
      - no *buffer\_head* exists, block not in memory
    - Aware of block information
      - *buffer\_head* exists, but block data (*b\_data*) not in memory
    - Aware of block information and data
      - Both the *buffer\_head* and its *b\_data* are valid ("\$ hit")

4

## Accessing blocks

- To read a block, FS uses *sb\_bread(...)*:
  - Find the corresponding *buffer\_head*
    - Create if doesn't exist
  - Make sure *buffer\_head->b\_data* is in memory (read from disk if necessary)
- To write a block:
  - *mark\_buffer\_dirty()* + *brelse()* - mark buffer as changed and release to kernel (which does the writing)

5

## Some buffer manager functions

<code>cse451_bread(pbh, inode, block, create)</code>	Get the <i>buffer_head</i> for the given disk block, ensuring that the data is in memory and ready for use. Increments ref count; always pair with a <i>brelse</i> .
<code>cse451_getblk(pbh, inode, block, create)</code>	Get the <i>buffer_head</i> for the given disk block. Does not guarantee anything about the state of the actual data. Increments ref count; always pair with a <i>brelse</i> . Zeros out new blocks (required for security).
<code>brelse(bh)</code>	Decrement the ref. count of the given buffer.
<code>mark_buffer_dirty(bh)</code>	Mark the buffer modified, meaning needs to be written to disk at some point.
<code>mark_buffer_uptodate(bh)</code>	Indicate that the data pointed to by bh is valid.

[ Remember this lock-release pattern for future use in multi-threaded (multi-process) programs; it's how reference-counted pointers also work. ]

6

## Network File Systems

- Provide access to remote files over a network
  - Typically aim for location and network **transparency**
- Designed and optimized for different types of operations
  - E.g., work over LANs, over WANs, support disconnected operations, fault-tolerance, scalability, consistency, etc.

Examples:

- Network File System (NFS) – Sun Microsystems
- Server Message Block (SMB) – originally IBM, Microsoft
- Andrew File System (AFS) – CMU
- Coda – CMU

7

## NFS

- A server exports (or shares) a directory
- A client mounts the remote directory onto his local FS namespace
  - The mounted directory looks like an integral subtree of the local file system, replacing the subtree descending from the local directory [1]
  - However, it's all namespace magic, nothing is actually stored on local disks

[1] <http://www.csie.fju.edu.tw/~yeh/courses/spring08/os/ch11.ppt>

8

## Mounting an NFS Export

- A remote exported directory can be "glued" onto the local hierarchy




Figure taken from <http://www.csie.fju.edu.tw/~yeh/courses/spring08/os/ch11.ppt>

9

## Mounting NFS Directories




Figure taken from <http://www.csie.fju.edu.tw/~yeh/courses/spring08/os/ch11.ppt>

10

## The NFS Protocol

- NFS is designed to operate in a heterogeneous environment of different machines, operating systems, and network architectures
  - NFS specifications are independent of these media
- This independence is achieved through the use of RPC and XDR (eXternal Data Representation)
- Nearly one-to-one correspondence between regular UNIX file system calls and the NFS protocol RPCs
  - looking up a file within a directory
  - reading a set of directory entries
  - accessing file attributes
  - reading and writing files

Bullets taken from presentation at <http://www.csie.fju.edu.tw/~yeh/courses/spring08/os/ch11.ppt>

11

## The NFS Architecture

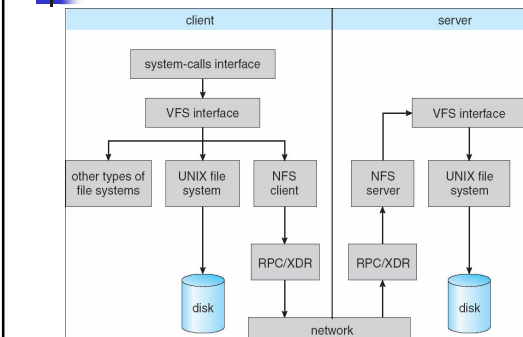


Figure taken from <http://www.csie.fju.edu.tw/~yeh/courses/spring08/os/ch11.ppt>

12



## Example: Setting up an NFS Share

- **Server exports directory**
  - Check `nfsd` is running and if not (e.g., `service nfsd start`)
  - Edit `/etc/exports` file
    - `/usr/shared 192.168.0.0/255.255.255.0(rw)`
    - `man exports` for more detailed structure of file
  - Force `nfsd` to re-read `/etc/exports` using `exportfs -ra`
- **Client mounts the remote directory locally**
  - `mount -t nfs 192.168.0.1:/usr/share /usr/local`  
(192.168.0.1 is the server's IP address)
  - can enable automatic mounting by editing `/etc/fstab` (man `fstab`)

**Note:**

- The above is just a "by-hand" example; use the Internet for more precise tutorials and troubleshooting

13