[Adapted for Fall 2025]

Using the page cache

Mitchell Gouzenko

Block Devices

- Filesystem are mounted on top of block devices
    - Ex: /dev/sda1
- Each device consists of contiguous data blocks.
- If you just try to read the raw bytes, it will look like garbage with bits of files scattered around.
    - But filesystems know how to treat the structure of data on disk, such that it makes sense to that specific FS.
    - Ex: ReiserFS will not be able to make sense out of a disk laid out in the format of ext4.
- You can write to device files, and the data WILL get written to the underlying disk.
    - If /dev/sdb is a USB drive, “echo foo > /dev/sdb” will write to that drive. `dd` utility often used for doing this.

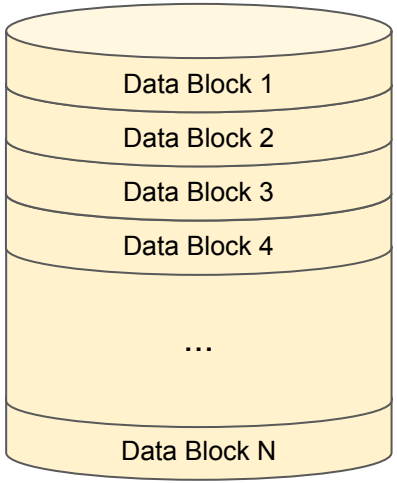

Block Devices (continued)

- When writing to regular file (for example, a file on an EXT4-formatted drive), data doesn't go to disk right away.
    - Note: device files like /dev/sda1 are not "regular" files.
- Data sits in kernel buffer, eventually will get synced to disk.
    - Advantage
        - Writing to buffer in kernel memory is much faster than writing to disk.
    - FS Buffer layer used to be called "Buffercache", but now the "Page Cache" (a different kernel subsystem) is used.
- Codepath for *actually* getting data to disk is lower-level than we need to worry about.
    - Involves disk scheduling and speaking to disk drivers (SATA II drivers, for instance)