

File System API

Creating a file

- `int fd = open("foo", O_CREAT|O_RDWR|O_TRUNC, S_IRUSR|S_IWUSR);`
- returns a file descriptor, a per-process integer that grants process a capability to perform certain operations
- `int close(int fd);` closes the file descriptor

Reading/Writing

- `ssize_t read (int fd, void *buf, size_t count);`
- `ssize_t write (int fd, void *buf, size_t count);`
 - return number of bytes read/written
- `off_t lseek (int fd, off_t offset, int whence);`
 - repositions file's offset (initially 0, updates on reads and writes)
 - to offset bytes (`SEEK_SET`)
 - to offset bytes from current offset (`SEEK_CUR`)
 - to offset bytes after the end of the file (`SEEK_END`)

File System API

Writing synchronously

- `int fsync (int fd);`
- flushes to disk all dirty data for file referred to by fd
- if file is newly created, must `fsync` also its directory!

Getting file's metadata

- `stat()`, `fstat()` – return a `stat` structure

```
struct stat {
    dev_t st_dev;      /* ID of device containing file */
    ino_t st_ino;      /* inode number */
    mode_t st_mode;    /* protection */
    nlink_t st_nlink;  /* number of hard links */
    uid_t st_uid;      /* user ID of owner */
    gid_t st_gid;      /* group ID of owner */
    dev_t st_rdev;     /* device ID (if special file) */
    off_t st_size;     /* total size, in bytes */
    blksize_t st_blksize; /* blocksize for filesystem I/O */
    blkcnt_t st_blocks; /* number of blocks allocated */
    time_t st_atime;   /* time of last access */
    time_t st_mtime;   /* time of last modification */
    time_t st_ctime;   /* time of last status change */
};
```

retrieved from file's inode

- on disk, per-file data structure
- may be cached in memory

File System API

Writing synchronously

- `int fsync (int fd);`
- flushes to disk all dirty data for file referred to by fd
- if file is newly created, must `fsync` also its directory!

Getting file's metadata

- `stat()`, `fstat()` – return a `stat` structure

```
struct stat {
    dev_t st_dev;      /* ID of device containing file */
    ino_t st_ino;      /* inode number */
    mode_t st_mode;    /* protection */
    nlink_t st_nlink;  /* number of hard links */
    uid_t st_uid;      /* user ID of owner */
    gid_t st_gid;      /* group ID of owner */
    dev_t st_rdev;     /* device ID (if special file) */
    off_t st_size;     /* total size, in bytes */
    blksize_t st_blksize; /* blocksize for filesystem I/O */
    blkcnt_t st_blocks; /* number of blocks allocated */
    time_t st_atime;   /* time of last access */
    time_t st_mtime;   /* time of last modification */
    time_t st_ctime;   /* time of last status change */
};
```

retrieved from file's inode

- on disk, per-file data structure
- may be cached in memory

Old Friends

Remember `fork()`?

```
int main(int argc, char *argv[]){
    int fd = open("file.txt", O_RDONLY);
    assert (fd >= 0);
    int rc = fork();
    if (rc == 0) {
        rc = lseek(fd, 10, SEEK_SET);
        printf("child: offset %d\n", rc);
    } else if (rc > 0) {
        (void) wait(NULL);
        printf("parent: offset %d\n",
              (int) lseek(fd, 10, SEEK_CUR));
    }
    return 0;
}
```

What does this code print?

Parent and child each have an independent file descriptor

Though independent, both correspond to the same integer (e.g., 3)

They both point to the same entry in the OS's Open File Table

An entry in that table looks like

```
struct file {
    int ref;
    char readable;
    char writable;
    struct inode *ip;
    uint off;
};
```

The reference count for file.txt would be 2!

The Directory

- ⦿ The directory holds instances of two types of mappings:
 - ▢ Hard links
 - ▶ map a file's human-friendly name (its local path) to the corresponding inode number
 - ▢ Symbolic (soft) links
 - ▶ maps a file's human-friendly name (its local path) to the number of an inode that contains the path name of a different file
 - you can think of it as a hard link for a special file, that indeed OS treats differently

Hard link No-Nos

- ⦿ Creating a hard link to a directory
 - ▢ may create a cycle in the directory tree!
- ⦿ Creating a hard link to files in other volumes
 - ▢ inode numbers are unique only within a single file system

Hard links

- ⦿ Creating file foo adds a hard link for file foo in the file's directory
 - ▢ `int link(const char *oldpath, const char *newpath)`
 - ▢ adds a hard link mapping path `newpath` to the inode number currently also mapped to file `oldpath`
 - ▢ invoked executing `ln` at the command line
- ⦿ Removing a file through the `rm [file]` command invokes a call to `int unlink(const char *pathname)`
 - ▢ removes from directory the hard link between `pathname` and corresponding inode number
- ⦿ Link count maintained in file's inode
 - ▢ inode reclaimed (file deleted) only when link count = 0; if file opened, wait to reclaim until file is closed

```
la13@en-cs-cisugcl10:~$ cd example
la13@en-cs-cisugcl10:~/example$ ls
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 ..
```

Example



```

la13@en-cs-cisugcl10:~$ cd example
la13@en-cs-cisugcl10:~/example$ ls
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 ..
la13@en-cs-cisugcl10:~/example$ echo ezra > cornell
la13@en-cs-cisugcl10:~/example$ cat cornell
ezra
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 .. 392852368 cornell

```

Example



```

la13@en-cs-cisugcl10:~$ cd example
la13@en-cs-cisugcl10:~/example$ ls
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 ..
la13@en-cs-cisugcl10:~/example$ echo ezra > cornell
la13@en-cs-cisugcl10:~/example$ cat cornell
ezra
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 .. 392852368 cornell
la13@en-cs-cisugcl10:~/example$ ln cornell bigred
la13@en-cs-cisugcl10:~/example$ cat bigred
ezra
la13@en-cs-cisugcl10:~/example$ ls -i
392852368 bigred 392852368 cornell

```

Example




```

la13@en-cs-cisugcl10:~$ cd example
la13@en-cs-cisugcl10:~/example$ ls
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 ..
la13@en-cs-cisugcl10:~/example$ echo ezra > cornell
la13@en-cs-cisugcl10:~/example$ cat cornell
ezra
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 .. 392852368 cornell
la13@en-cs-cisugcl10:~/example$ ln cornell bigred
la13@en-cs-cisugcl10:~/example$ cat bigred
ezra
la13@en-cs-cisugcl10:~/example$ ls -i
392852368 bigred 392852368 cornell
la13@en-cs-cisugcl10:~/example$ ln bigred ../bestivy
la13@en-cs-cisugcl10:~/example$ ls -i
392852368 bigred 392852368 cornell
la13@en-cs-cisugcl10:~/example$ cd ..
la13@en-cs-cisugcl10:~$ cat bestivy
ezra
la13@en-cs-cisugcl10:~$ ls -i
392852368 bestivy 398842589 CS4410-2020sp-A4 392852366 example

```

Example



```

la13@en-cs-cisugcl10:~$ cd example
la13@en-cs-cisugcl10:~/example$ ls
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 ..
la13@en-cs-cisugcl10:~/example$ echo ezra > cornell
la13@en-cs-cisugcl10:~/example$ cat cornell
ezra
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 .. 392852368 cornell
la13@en-cs-cisugcl10:~/example$ ln cornell bigred
la13@en-cs-cisugcl10:~/example$ cat bigred
ezra
la13@en-cs-cisugcl10:~/example$ ls -i
392852368 bigred 392852368 cornell
la13@en-cs-cisugcl10:~/example$ ln bigred ../bestivy
la13@en-cs-cisugcl10:~/example$ ls -i
392852368 bigred 392852368 cornell
la13@en-cs-cisugcl10:~/example$ cd ..
la13@en-cs-cisugcl10:~$ cat bestivy
ezra
la13@en-cs-cisugcl10:~$ ls -i
392852368 bestivy 398842589 CS4410-2020sp-A4 392852366 example
la13@en-cs-cisugcl10:~$ cd example
la13@en-cs-cisugcl10:~/example$ rm cornell
la13@en-cs-cisugcl10:~/example$ rm bigred
la13@en-cs-cisugcl10:~/example$ ls -i
la13@en-cs-cisugcl10:~/example$ cd ..

```

Example

Example

```
la13@en-cs-cisugcl10:~$ cd example
la13@en-cs-cisugcl10:~/example$ ls
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 ..
la13@en-cs-cisugcl10:~/example$ echo ezra > cornell
la13@en-cs-cisugcl10:~/example$ cat cornell
ezra
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 .. 392852368 cornell
la13@en-cs-cisugcl10:~/example$ ln cornell bigred
la13@en-cs-cisugcl10:~/example$ cat bigred
ezra
la13@en-cs-cisugcl10:~/example$ ls -i
392852368 bigred 392852368 cornell
la13@en-cs-cisugcl10:~/example$ ln bigred ../bestivy
la13@en-cs-cisugcl10:~/example$ ls -i
392852368 bigred 392852368 cornell
la13@en-cs-cisugcl10:~/example$ cd ..
la13@en-cs-cisugcl10:~$ cat bestivy
ezra
la13@en-cs-cisugcl10:~$ ls -i
392852368 bestivy 398842589 CS4410-2020sp-A4 392852366 example
```

Example

```
la13@en-cs-cisugcl10:~$ cd example
la13@en-cs-cisugcl10:~/example$ ls
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 ..
la13@en-cs-cisugcl10:~/example$ echo ezra > cornell
la13@en-cs-cisugcl10:~/example$ cat cornell
ezra
la13@en-cs-cisugcl10:~/example$ ls -ai
392852366 . 391230414 .. 392852368 cornell
la13@en-cs-cisugcl10:~/example$ ln cornell bigred
la13@en-cs-cisugcl10:~/example$ cat bigred
ezra
la13@en-cs-cisugcl10:~/example$ ls -i
392852368 bigred 392852368 cornell
la13@en-cs-cisugcl10:~/example$ ln bigred ../bestivy
la13@en-cs-cisugcl10:~/example$ ls -i
392852368 bigred 392852368 cornell
la13@en-cs-cisugcl10:~/example$ cd ..
la13@en-cs-cisugcl10:~$ cat bestivy
ezra
la13@en-cs-cisugcl10:~$ ls -i
392852368 bestivy 398842589 CS4410-2020sp-A4 392852366 example
la13@en-cs-cisugcl10:~$
```

Symbolic (Soft) links

- More flexible than hard links
 - can link to a directory
 - can link to files in another volume
- A map between pathnames
 - to link newpathname to existingpathname for file inode1:
 - ▶ create a hard link between newpathname and new file inode2
 - ▶ store in inode2 the existingpathname for inode1
 - so, a symbolic link is really a file (inode2 in our example) of a third type
 - ▶ neither a regular file nor a directory
- Created using ln, but with the -s flag

Example

Example

~/example/cornell



```

lail3@en-cs-cisugc105:~$ cd example
lail3@en-cs-cisugc105:~/example$ echo ezra > cornell
lail3@en-cs-cisugc105:~/example$ ls -l
392852367 cornell

```

Example

~/example/cornell



~/example/bigred

```

lail3@en-cs-cisugc105:~$ cd example
lail3@en-cs-cisugc105:~/example$ echo ezra > cornell
lail3@en-cs-cisugc105:~/example$ ls -l
392852367 cornell
lail3@en-cs-cisugc105:~/example$ ln cornell bigred
lail3@en-cs-cisugc105:~/example$ ls -l
392852367 bigred 392852367 cornell

```

Example

~/highabove



~/bestivvy

~/example/bigred

```

lail3@en-cs-cisugc105:~$ cd example
lail3@en-cs-cisugc105:~/example$ echo ezra > cornell
lail3@en-cs-cisugc105:~/example$ ls -l
392852367 cornell
lail3@en-cs-cisugc105:~/example$ ln cornell bigred
lail3@en-cs-cisugc105:~/example$ ls -l
392852367 bigred 392852367 cornell
lail3@en-cs-cisugc105:~/example$ cd ..
lail3@en-cs-cisugc105:~$ ln example/cornell bestivvy
lail3@en-cs-cisugc105:~$ ln -s example/cornell highabove
lail3@en-cs-cisugc105:~$ ls -l
392852367 bestivvy 398842589 CS4410-2020sp-A4 392852366 example 392971138 highabove

```

Example

~/highabove



~/bestivvy

~/example/bigred


```

lail3@en-cs-cisugc105:~$ cd example
lail3@en-cs-cisugc105:~/example$ echo ezra > cornell
lail3@en-cs-cisugc105:~/example$ ls -l
392852367 cornell
lail3@en-cs-cisugc105:~/example$ ln cornell bigred
lail3@en-cs-cisugc105:~/example$ ls -l
392852367 bigred 392852367 cornell
lail3@en-cs-cisugc105:~/example$ cd ..
lail3@en-cs-cisugc105:~$ ln example/cornell bestivvy
lail3@en-cs-cisugc105:~$ ln -s example/cornell highabove
lail3@en-cs-cisugc105:~$ ls -l
392852367 bestivvy 398842589 CS4410-2020sp-A4 392852366 example 392971138 highabove
lail3@en-cs-cisugc105:~$ ls -l
total 8
-rw-r--r-- 3 lail3 pug-lal3 5 Apr 28 23:03 bestivvy
drwxr-xr-x 4 lail3 pug-lal3 4096 Apr 27 11:55 CS4410-2020sp-A4
drwxr-xr-x 2 lail3 pug-lal3 4096 Apr 28 23:03 example
lrwxrwxrwx 1 lail3 pug-lal3 15 Apr 28 23:04 highabove -> example/cornell

```

Example

~/highabove



~/bestivty


```

la13@en-cs-cisugcl05:~$ cd example
la13@en-cs-cisugcl05:~/example$ echo ezra > cornell
la13@en-cs-cisugcl05:~/example$ ls -l
392852367 cornell
la13@en-cs-cisugcl05:~/example$ ln cornell bigred
la13@en-cs-cisugcl05:~/example$ ls -l
392852367 bigred 392852367 cornell
la13@en-cs-cisugcl05:~/example$ cd ..
la13@en-cs-cisugcl05:~$ ln example/cornell bestivty
la13@en-cs-cisugcl05:~$ ln -s example/cornell highabove
la13@en-cs-cisugcl05:~$ ls -l
392852367 bestivty 398842589 CS4410-2020sp-A4 392852366 example 392971138 highabove
la13@en-cs-cisugcl05:~$ ls -l
total 8
-rw-r--r-- 3 la13 pug-lal3 5 Apr 28 23:03 bestivty
drwxr-xr-x 4 la13 pug-lal3 4096 Apr 27 11:55 CS4410-2020sp-A4
drwxr-xr-x 2 la13 pug-lal3 4096 Apr 28 23:03 example
lrwxrwxrwx 1 la13 pug-lal3 15 Apr 28 23:04 highabove -> example/cornell
la13@en-cs-cisugcl05:~$ cat bestivty
ezra
la13@en-cs-cisugcl05:~$ cat highabove
ezra

```

Example

~/highabove



~/bestivty


```

la13@en-cs-cisugcl05:~$ cd example
la13@en-cs-cisugcl05:~/example$ echo ezra > cornell
la13@en-cs-cisugcl05:~/example$ ls -l
392852367 cornell
la13@en-cs-cisugcl05:~/example$ ln cornell bigred
la13@en-cs-cisugcl05:~/example$ ls -l
392852367 bigred 392852367 cornell
la13@en-cs-cisugcl05:~/example$ cd ..
la13@en-cs-cisugcl05:~$ ln example/cornell bestivty
la13@en-cs-cisugcl05:~$ ln -s example/cornell highabove
la13@en-cs-cisugcl05:~$ ls -l
392852367 bestivty 398842589 CS4410-2020sp-A4 392852366 example 392971138 highabove
la13@en-cs-cisugcl05:~$ ls -l
total 8
-rw-r--r-- 3 la13 pug-lal3 5 Apr 28 23:03 bestivty
drwxr-xr-x 4 la13 pug-lal3 4096 Apr 27 11:55 CS4410-2020sp-A4
drwxr-xr-x 2 la13 pug-lal3 4096 Apr 28 23:03 example
lrwxrwxrwx 1 la13 pug-lal3 15 Apr 28 23:04 highabove -> example/cornell
la13@en-cs-cisugcl05:~$ cat bestivty
ezra
la13@en-cs-cisugcl05:~$ cat highabove
ezra
la13@en-cs-cisugcl05:~$ rm example/cornell

```

Example

~/highabove



~/bestivty


```

la13@en-cs-cisugcl05:~$ cd example
la13@en-cs-cisugcl05:~/example$ echo ezra > cornell
la13@en-cs-cisugcl05:~/example$ ls -l
392852367 cornell
la13@en-cs-cisugcl05:~/example$ ln cornell bigred
la13@en-cs-cisugcl05:~/example$ ls -l
392852367 bigred 392852367 cornell
la13@en-cs-cisugcl05:~/example$ cd ..
la13@en-cs-cisugcl05:~$ ln example/cornell bestivty
la13@en-cs-cisugcl05:~$ ln -s example/cornell highabove
la13@en-cs-cisugcl05:~$ ls -l
392852367 bestivty 398842589 CS4410-2020sp-A4 392852366 example 392971138 highabove
la13@en-cs-cisugcl05:~$ ls -l
total 8
-rw-r--r-- 3 la13 pug-lal3 5 Apr 28 23:03 bestivty
drwxr-xr-x 4 la13 pug-lal3 4096 Apr 27 11:55 CS4410-2020sp-A4
drwxr-xr-x 2 la13 pug-lal3 4096 Apr 28 23:03 example
lrwxrwxrwx 1 la13 pug-lal3 15 Apr 28 23:04 highabove -> example/cornell
la13@en-cs-cisugcl05:~$ cat bestivty
ezra
la13@en-cs-cisugcl05:~$ cat highabove
ezra
la13@en-cs-cisugcl05:~$ rm example/cornell
la13@en-cs-cisugcl05:~$ cat bestivty
ezra

```

Example

~/highabove



~/bestivty

```

la13@en-cs-cisugcl05:~$ cd example
la13@en-cs-cisugcl05:~/example$ echo ezra > cornell
la13@en-cs-cisugcl05:~/example$ ls -l
392852367 cornell
la13@en-cs-cisugcl05:~/example$ ln cornell bigred
la13@en-cs-cisugcl05:~/example$ ls -l
392852367 bigred 392852367 cornell
la13@en-cs-cisugcl05:~/example$ cd ..
la13@en-cs-cisugcl05:~$ ln example/cornell bestivty
la13@en-cs-cisugcl05:~$ ln -s example/cornell highabove
la13@en-cs-cisugcl05:~$ ls -l
392852367 bestivty 398842589 CS4410-2020sp-A4 392852366 example 392971138 highabove
la13@en-cs-cisugcl05:~$ ls -l
total 8
-rw-r--r-- 3 la13 pug-lal3 5 Apr 28 23:03 bestivty
drwxr-xr-x 4 la13 pug-lal3 4096 Apr 27 11:55 CS4410-2020sp-A4
drwxr-xr-x 2 la13 pug-lal3 4096 Apr 28 23:03 example
lrwxrwxrwx 1 la13 pug-lal3 15 Apr 28 23:04 highabove -> example/cornell
la13@en-cs-cisugcl05:~$ cat bestivty
ezra
la13@en-cs-cisugcl05:~$ cat highabove
ezra
la13@en-cs-cisugcl05:~$ rm example/cornell
la13@en-cs-cisugcl05:~$ cat bestivty
ezra
la13@en-cs-cisugcl05:~$ cat highabove
cat: highabove: No such file or directory
la13@en-cs-cisugcl05:~$

```

Permission Bits

```
lal3@en-es-cisugc105:~$ ls -l
total 8
-rw-r--r-- 3 lal3 pug-lal3  5 Apr 28 23:03 bestivy
drwxr-sr-x 4 lal3 pug-lal3 4096 Apr 27 11:55 CS4410-2020sp-A4
drwxr-sr-x 2 lal3 pug-lal3 4096 Apr 28 23:03 example
lrwxrwxrwx 1 lal3 pug-lal3  15 Apr 28 23:04 highabove -> example/cornell
```

File bestivy

- leading - says bestivy is a regular file
 - d is for directory; l is for soft link
- Next nine characters are permission bits
 - rwX for owner, group, everyone
 - owner can read and write; group and others can just read
 - x set in a regular file means means file can be executed
 - x set in a directory that user/group/everybody is allow to cd to that directory
 - can be set using chmod

Mount

- Mount: allows multiple file systems on multiple volumes to form a single logical hierarchy
 - a mapping from some path in existing file system to the root directory of the mounted file system

