

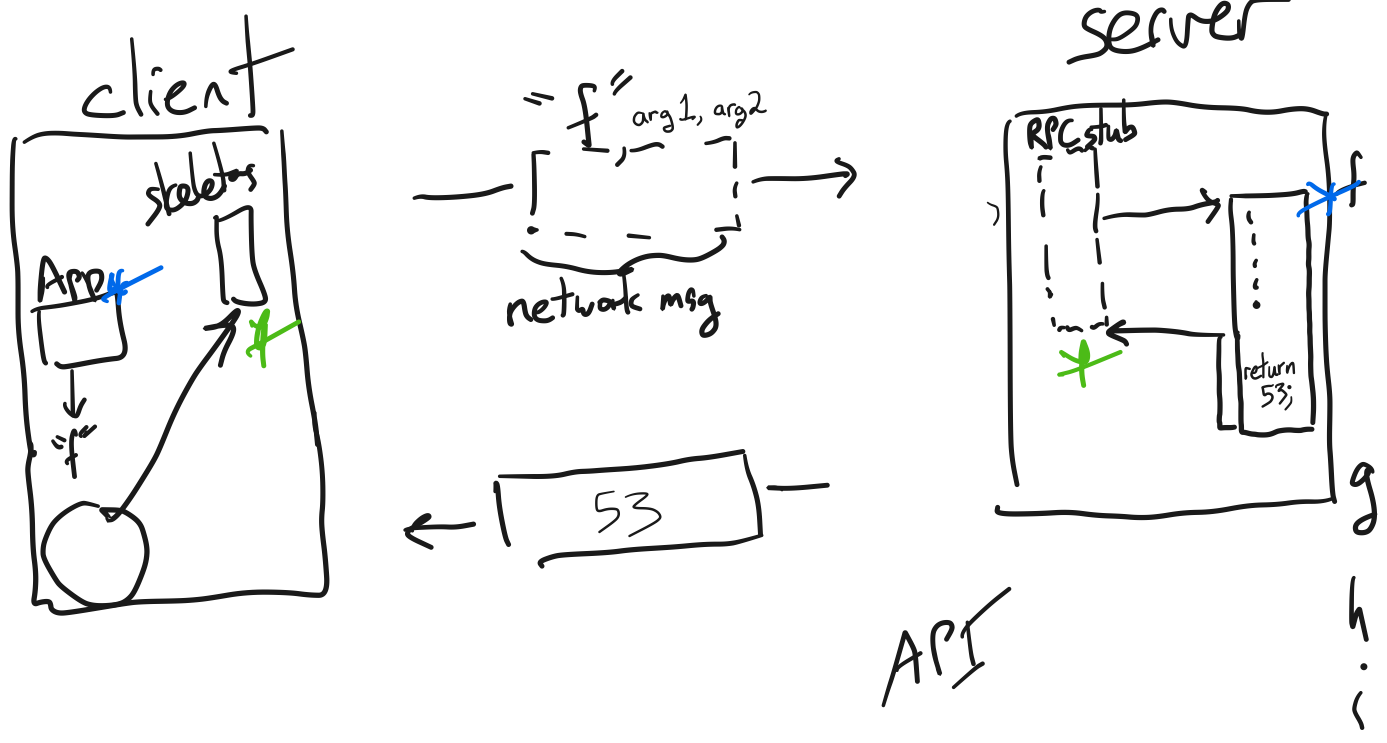


- 1. Last time
- 2. Crash recovery, continued
- 3. RPC, client/server systems
- 4. NFS (case study of RPC + client/server)
 - Intro + background
 - How it works
 - lab5 detour
 - Statelessness (protocol statelessness)
 - Transparency



redo+undo logging

3. RPC, client server

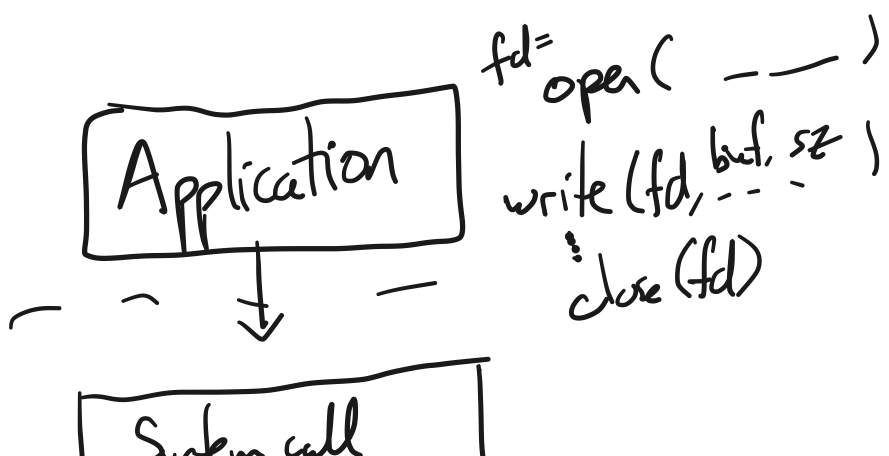




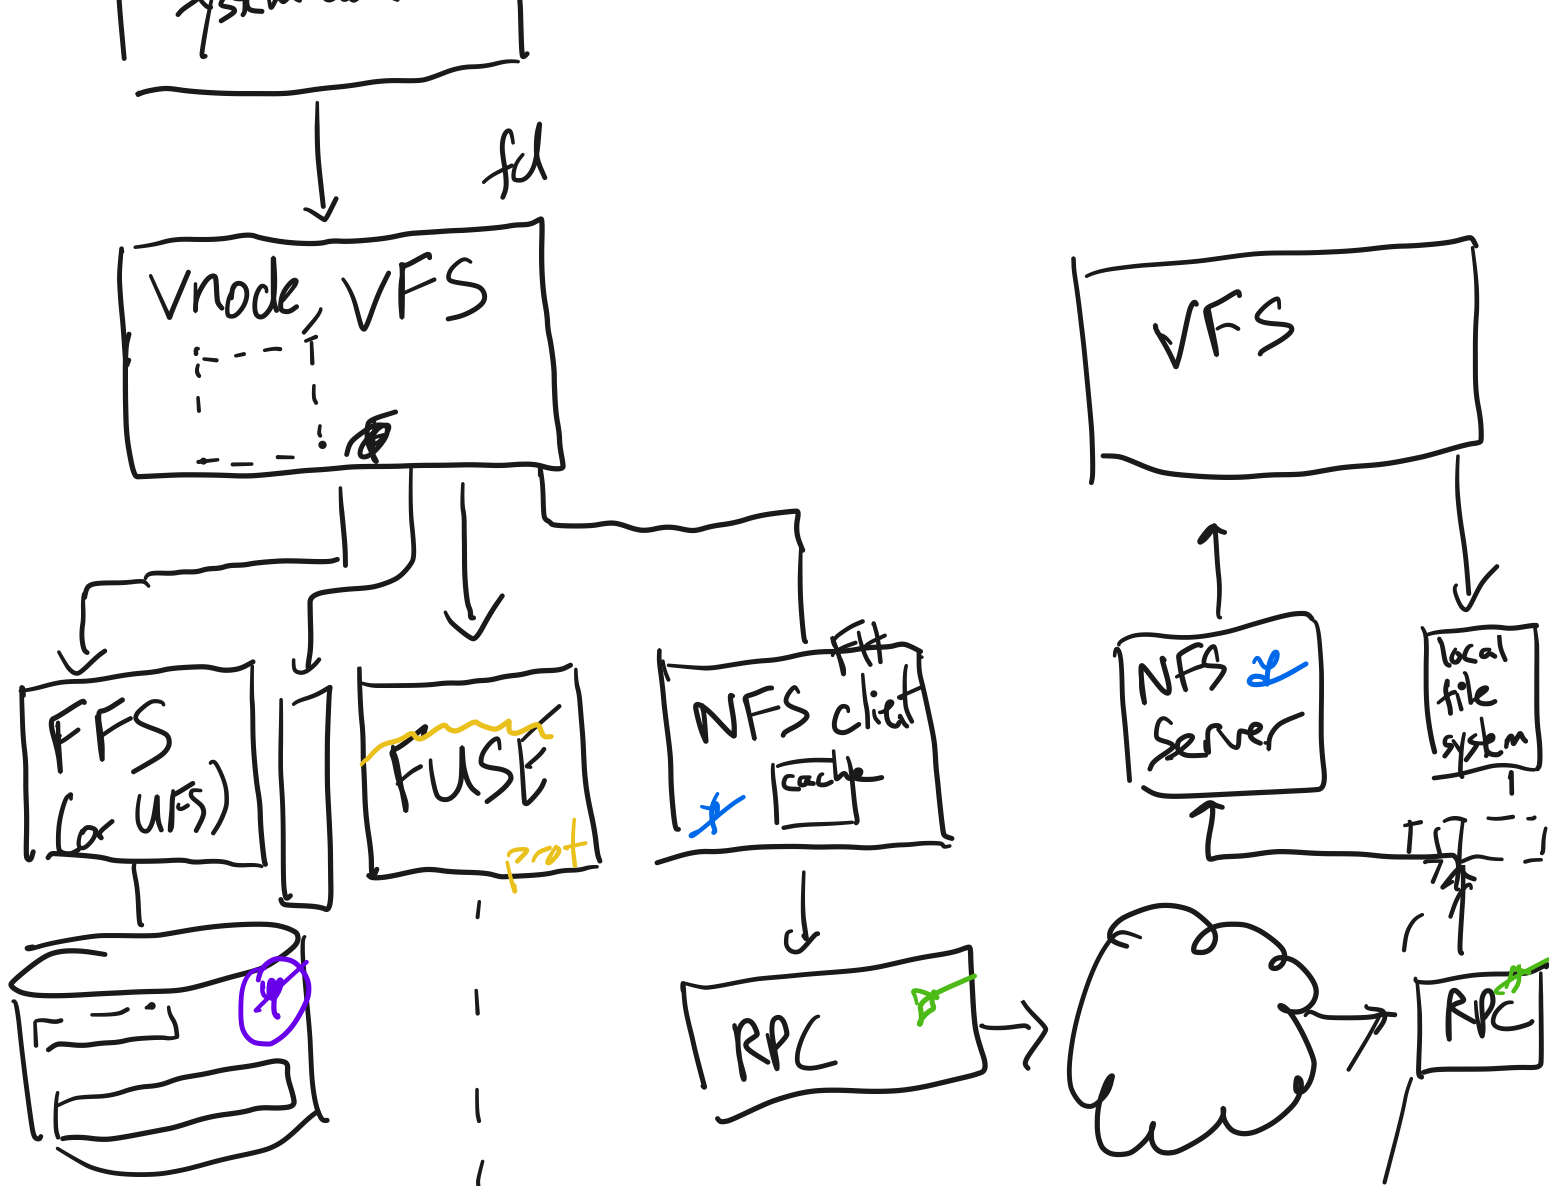
4. NFS

\$ ls

Client

Server






lab notes:
 → be aware: each inode gets its
 own disk block
 → pointer-to-pointer
 func...#ppblocks, ...)

f (char ...)
 info ...

Client

```
open("/usr/jo/lab.c", ...)
/usr/jo/lab.c, WR_ONLY
```



Server

```
LOOKUP(fh0, "/usr")
```

file

```
fh1 = [FS id | inode# | generation#]
```

```
LOOKUP(fh1, "/usr/jo")
```

```
LOOKUP(fh2, "/usr/jo/lab.c")
```

```
fh = [FS id | inode# | generation#]
```

could be as a result of close()

```
WRITE(fh, position, data, size)
```

return code

WRITE: idempotent

RENAMING (a, b)

READ :

rename :

?

~~CREATE~~
~~RENAME ("a", "b")~~

← error

CREATE("f" ...)

← ok


~~CREATE("f" ...)~~

← ERROR

~~→~~

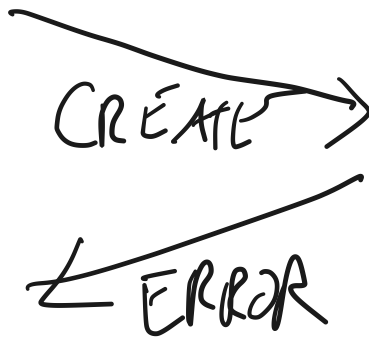

?

←



4. Transparency (we don't have it always!)

- One issue: gen #s added to legacy FS.
- Successful ops can return errors



↓ flushes

..."

✓ "OK

caching

read-caching

write-caching

A: write();
/* caching */

B: open();
read();
close();

A
write();
/* caching */

close();

B
open();
read();

if ((rc = close(fd)) < 0)
error - - -

no guarantee
the data is
seen.