

OSTEP Chapter 42


ECE 3600, Fall 2022

Table of Contents


- [1. Crash Consistency](#)
- [2. fsck](#)
- [3. Data Journaling](#)
- [4. Exercises](#)

1. Crash Consistency

Initial file system, one file, one block: $B[v1]$, $I[v1]$, D_a



Data block D_b appended, requires 3 writes: $B[v2]$, $I[v2]$, D_b



Crash Scenarios:

1. $B[v2]$ written, but not $I[v2]$ and D_b
2. $I[v2]$ written, but not $B[v2]$ and D_b
3. D_b written, but not $B[v2]$ and $I[v2]$
4. $B[v2]$ and $I[v2]$ written, but not D_b
5. $I[v2]$ and D_b written, but not $B[v2]$
6. $B[v2]$ and D_b written, but not $I[v2]$

Which leave file system inconsistent?

How to fix?

2. fsck

File System Checks:


- Superblock
- Free blocks: inode and data bitmaps
- Inode state
- Inode link counts
- Duplicates
- Bad blocks
- Directory consistency

Slow.


3. Data Journaling

Write-ahead logging


Without journal:



With journal:



Journal:



4. Exercises

Using [fsck.py](#):

```
$ python ./fsck.py -n 6 -i 8 -d 8 -c
```

Initial state of file system:

```
          01234567
inode bitmap 10001101

inodes      0          1 2 3 4          5          6 7
            [d a:0 r:3] [] [] [] [d a:3 r:2] [f a:-1 r:1] [] [f a:6 r:1]

          01234567
data bitmap 10010010

data        0          1 2 3          4 5 6 7
            [(.,0) (.,0) (g,4) (t,7)] [] [] [(.,4) (.,0) (x,5)] [] [] [z] []
```

--> determine which files and directories are there

Final state of file system:

```
          01234567
inode bitmap 10001101

inodes      0          1 2 3 4          5          6 7
            [d a:0 r:3] [] [] [] [d a:4 r:2] [f a:-1 r:1] [] [f a:6 r:1]

          01234567
data bitmap 10010010

data        0          1 2 3          4 5 6 7
            [(.,0) (.,0) (g,4) (t,7)] [] [] [(.,4) (.,0) (x,5)] [] [] [z] []
```

--> CORRUPTION = ?