

L15: Transactions

Frans Kaashoek & Nickolai Zeldovich
6.033 Spring 2012

Bank account transfer

xfer(bank, a, b, amt):

bank[a] = bank[a] - amt

bank[b] = bank[b] + amt

Bank account transfer

xfer(bank, a, b, 50):

xfer(bank, a, b, amt):

bank[a] = bank[a] - amt

bank[b] = bank[b] + amt

← *a=100, b=100*

← *a=50, b=100*

← *a=50, b=150*

Bank account transfer

```
xfer(bank, a, b, amt):  
    bank[a] = bank[a] - amt  
    bank[b] = bank[b] + amt
```

```
audit(bank):  
    sum = 0  
    for acct in bank:  
        sum = sum + bank[acct]  
    return sum
```

Bank account transfer

audit(bank):

xfer(bank, a, b, amt):

bank[a] = bank[a] - amt

bank[b] = bank[b] + amt

← *sum=200*

← *sum=150*

← *sum=200*

audit(bank):

sum = 0

for acct in bank:

sum = sum + bank[acct]

return sum

Eventual goal: transactions all-or-nothing & before-or-after atomicity

xfer(bank, a, b, amt):

begin

bank[a] = bank[a] - amt

bank[b] = bank[b] + amt

commit

audit(bank):

begin

sum = 0

for acct in bank:

 sum = sum + bank[acct]

commit

return sum

Strawman implementation

```
xfer(bankfile, a, b, amt):  
    bank = read_accounts(bankfile)  
    bank[a] = bank[a] - amt  
    bank[b] = bank[b] + amt  
    write_accounts(bankfile)
```

Shadow copy

```
xfer(bankfile, a, b, amt):  
    bank = read_accounts(bankfile)  
    bank[a] = bank[a] - amt  
    bank[b] = bank[b] + amt  
    write_accounts("#bankfile")  
    rename("#bankfile", bankfile)
```


File system data structures

directory data blocks:

filename "bank" → inode 12

filename "#bank" → inode 13

inode 12:

data blocks: 3, 4, 5

refcount: 1

inode 13:

data blocks: 6, 7, 8

refcount: 1

rename("#bank", "bank")

directory data blocks:

filename "bank" → inode 12

filename "#bank" → inode 13

inode 12:

data blocks: 3, 4, 5

refcount: 1

inode 13:

data blocks: 6, 7, 8

refcount: 1

rename("#bank", "bank")

directory data blocks:

filename "bank" → inode 12

filename "#bank" → inode 13

inode 12:

data blocks: 3, 4, 5

refcount: 1

inode 13:

data blocks: 6, 7, 8

refcount: **2**

rename("#bank", "bank")

directory data blocks:

filename "bank" → inode **13**

filename "#bank" → inode 13

inode 12:

data blocks: 3, 4, 5

refcount: 1

inode 13:

data blocks: 6, 7, 8

refcount: **2**

rename("#bank", "bank")

directory data blocks:

filename "bank" → inode **13**

filename "#bank" → inode 13

inode 12:

data blocks: 3, 4, 5

refcount: **0**

inode 13:

data blocks: 6, 7, 8

refcount: **2**

rename("#bank", "bank")

directory data blocks:

filename "bank" → inode **13**

~~filename "#bank" ⇒ inode 13~~

inode 12:

data blocks: 3, 4, 5

refcount: **0**

inode 13:

data blocks: 6, 7, 8

refcount: **2**

rename("#bank", "bank")

directory data blocks:

filename "bank" → inode **13**

~~filename "#bank" ⇒ inode 13~~

inode 12:

data blocks: 3, 4, 5

refcount: **0**

inode 13:

data blocks: 6, 7, 8

refcount: **1**

Recovery after crash

```
salvage(disk):  
    for inode in disk.inodes:  
        inode.refcnt =  
            find_all_refs(disk.root_dir, inode)  
  
    if exists("#bank"):  
        unlink("#bank")
```