Autopsy of a multiserver deadlock in the HelenOS filesystem layer

Jakub Jermář

Introduction

Microkernel + Multiserver

=

lots of message passing among lots of processes

=

breeding ground for distributed deadlocks

The HelenOS usecase

Hang on to your hats as we will go back in time to 2011.

Setting a collision course

- HelenOS mainline,1219

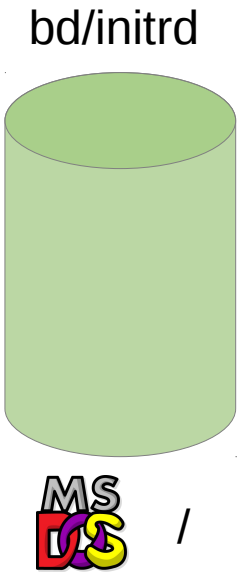

Setting a collision course

- Create a sufficiently large file

```
# / mkfile --size 300k img
```

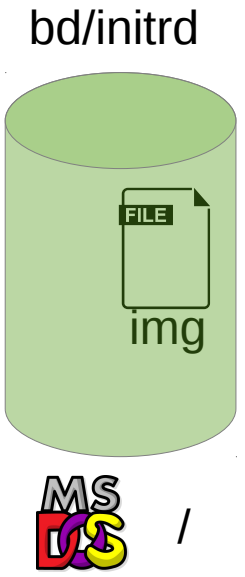

Setting a collision course

- Create a file-backed block device on top of it

```
# / file_bd img loop
```

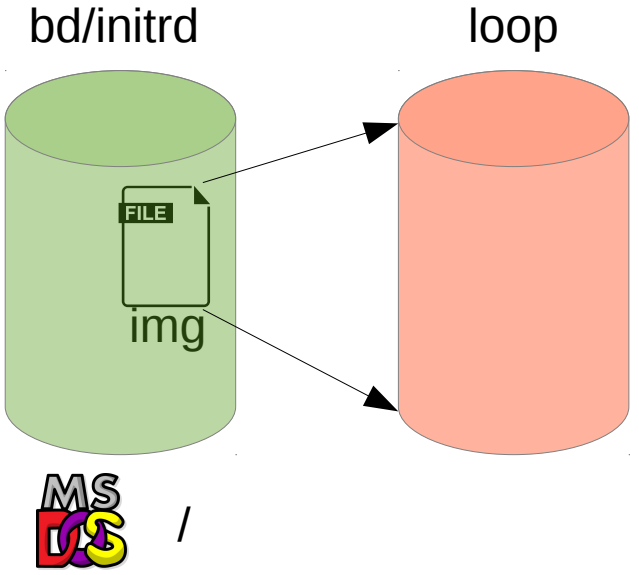

Setting a collision course

- Format as a MINIX file system

```
# / mkmfs loop
```

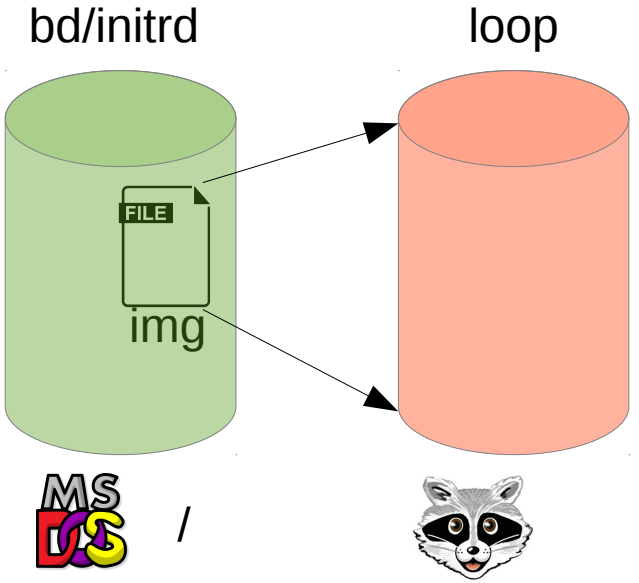

Setting a collision course

- And try to mount it under /data

```
# / mount mfs /data loop
```

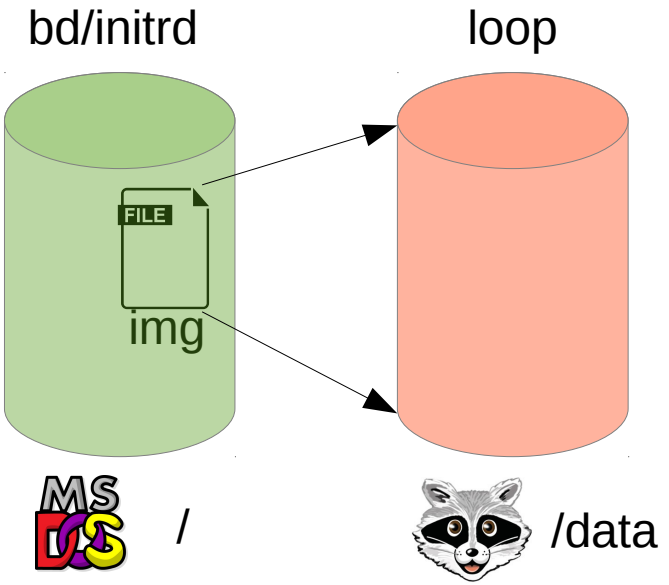

Setting a collision course

- And try to mount it under /data

```
# / mount mfs /data loop
```

- …it will not return

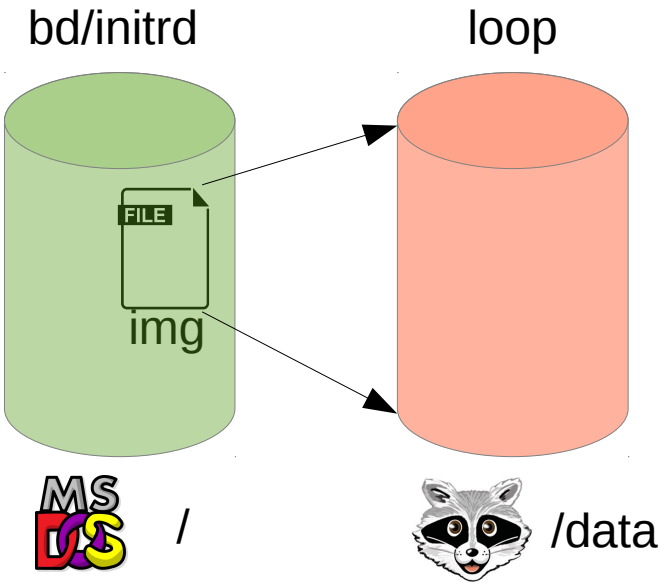


```
HelenOS release 0.4.3 (Sashimi), revision 1219M (martin@decky.cz-20110909154621-97c015mxwdo1q8b2)
Built on 2015-01-23 21:44:36
Running on amd64 (/loc/term/vc0)
Copyright (c) 2001-2011 HelenOS project

Welcome to HelenOS!
http://www.helenos.org/

Type 'help' [Enter] to see a few survival tips.

/ # mkfile --size 300k img
/ # file_bd img loop
file_bd: File-backed block device driver
file_bd: Accepting connections
/ # mkmfs loop
mkmfs: Block device has 600 blocks.
mkmfs: Creating Minix file system on device
mkmfs: 4096 block size
mkmfs: 128 inodes
mkmfs: 75 zones
mkmfs: inode table blocks = 2
mkmfs: inode bitmap blocks = 1
mkmfs: zone bitmap blocks = 1
mkmfs: first data zone = 6
mkmfs: max file size = 2147483647
mkmfs: long fnames = No
/ # mfs
mfs: HelenOS Minix file system server
mfs: Accepting connections
/ # mount mfs /data loop
```

Analysis

- Perhaps we could try to use *kconsole* to investigate

```
# / kcon
kconsole>
```

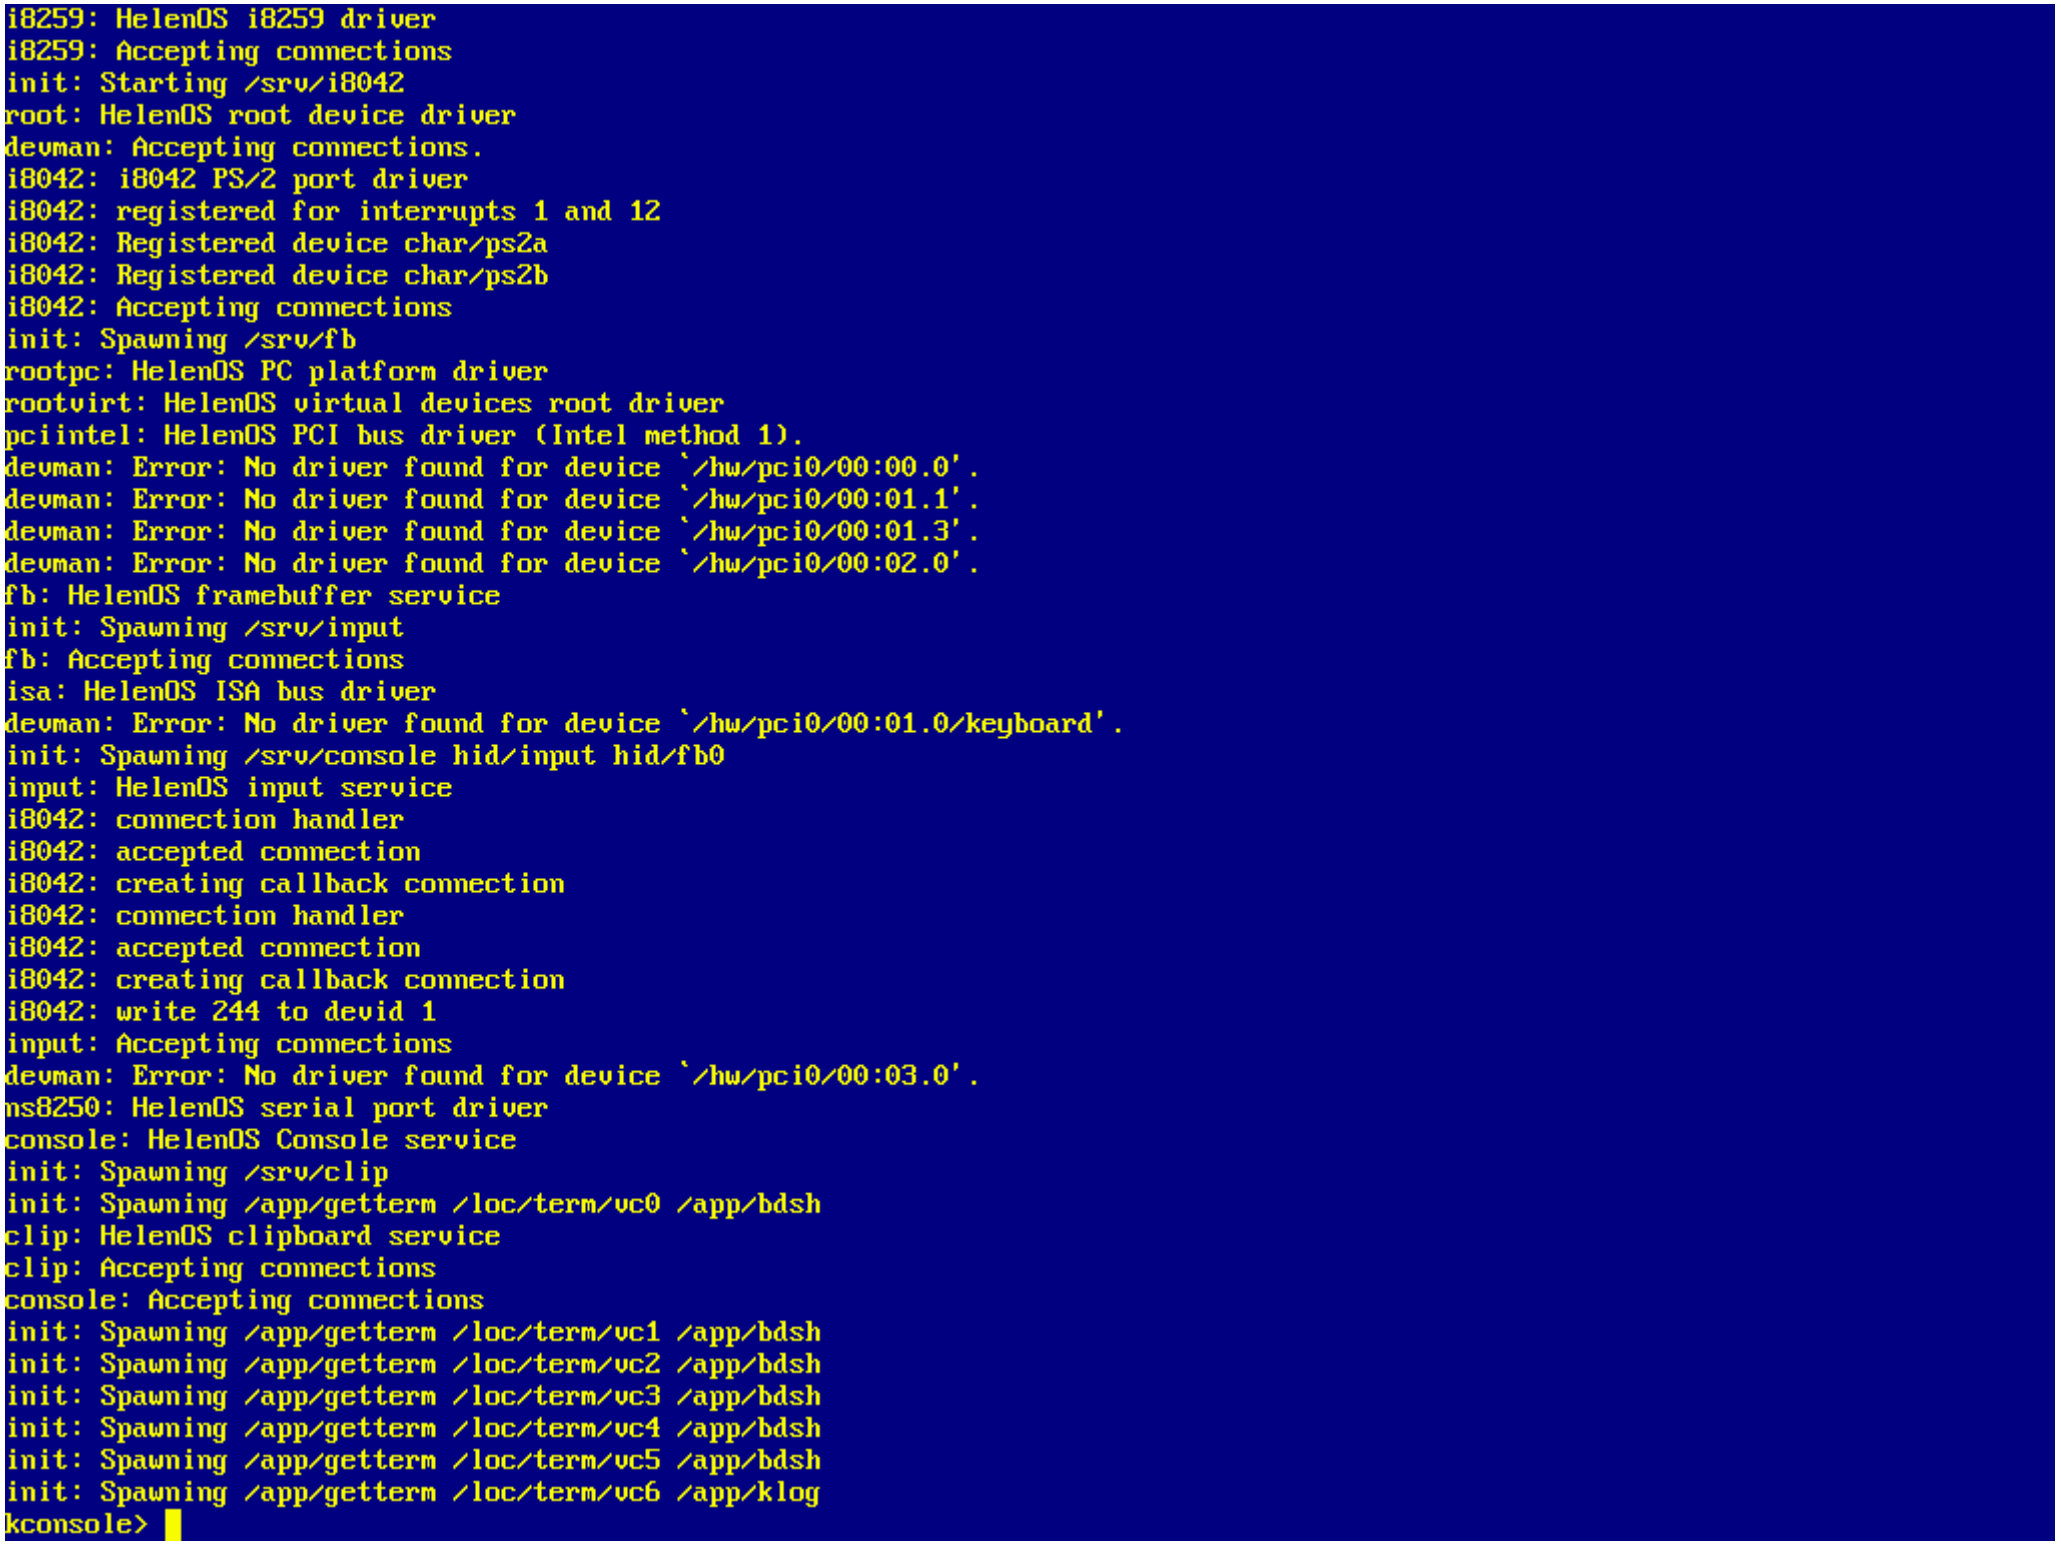
i8259: HelenOS i8259 driver
i8259: Accepting connections
init: Starting /srv/i8042
root: HelenOS root device driver
devman: Accepting connections.
i8042: i8042 PS/2 port driver
i8042: registered for interrupts 1 and 12
i8042: Registered device char/ps2a
i8042: Registered device char/ps2b
i8042: Accepting connections
init: Spawning /srv/fb
rootpc: HelenOS PC platform driver
rootvirt: HelenOS virtual devices root driver
pciintel: HelenOS PCI bus driver (Intel method 1).
devman: Error: No driver found for device `/hw/pci0/00:00.0'.
devman: Error: No driver found for device `/hw/pci0/00:01.1'.
devman: Error: No driver found for device `/hw/pci0/00:01.3'.
devman: Error: No driver found for device `/hw/pci0/00:02.0'.
fb: HelenOS framebuffer service
init: Spawning /srv/input
fb: Accepting connections
isa: HelenOS ISA bus driver
devman: Error: No driver found for device `/hw/pci0/00:01.0/keyboard'.
init: Spawning /srv/console hid/input hid/fb0
input: HelenOS input service
i8042: connection handler
i8042: accepted connection
i8042: creating callback connection
i8042: connection handler
i8042: accepted connection
i8042: creating callback connection
i8042: write 244 to devid 1
input: Accepting connections
devman: Error: No driver found for device `/hw/pci0/00:03.0'.
ns8250: HelenOS serial port driver
console: HelenOS Console service
init: Spawning /srv/clip
init: Spawning /app/getterm /loc/term/vc0 /app/bdsh
clip: HelenOS clipboard service
clip: Accepting connections
console: Accepting connections
init: Spawning /app/getterm /loc/term/vc1 /app/bdsh
init: Spawning /app/getterm /loc/term/vc2 /app/bdsh
init: Spawning /app/getterm /loc/term/vc3 /app/bdsh
init: Spawning /app/getterm /loc/term/vc4 /app/bdsh
init: Spawning /app/getterm /loc/term/vc5 /app/bdsh
init: Spawning /app/getterm /loc/term/vc6 /app/klog
kconsole>

Analysis

- List all processes

```
kconsole> tasks
[id    ] [name     ] [ctn] [address          ] [as               ]
1         kernel       0     0xffff80000004a000 0xffff80000009c000
2         init:ns      0     0xffff80000005c000 0xffff80000009c088
4         init:loc     0     0xffff800000068000 0xffff80000009c198
5         init:rd      0     0xffff80000006c000 0xffff80000009c220
6         init:vfs     0     0xffff800000074000 0xffff80000009c2a8
7         init:fat     0     0xffff80000007a000 0xffff80000009c330
8         tmpfs        0     0xffff800007f36000 0xffff80000009c3b8
9         locfs        0     0xffff800007f76000 0xffff80000009c440
…
26        getterm      0     0xffff800007ab2000 0xffff80000009ccc0
27        bdsh         0     0xffff800007ade000 0xffff80000009cd48
…
38        klog         0     0xffff800000062000 0xffff80000009c110
39        file_bd      0     0xffff800007120000 0xffff800007bc12a8
41        mfs          0     0xffff8000071fc000 0xffff800007bc13b8
```

Analysis

- List all processes

```
kconsole> tasks
[id    ] [name     ] [ctn] [address          ] [as               ]
1         kernel       0     0xffff80000004a000 0xffff80000009c000
2         init:ns      0     0xffff80000005c000 0xffff80000009c088
4         init:loc     0     0xffff800000068000 0xffff80000009c198
5         init:rd      0     0xffff80000006c000 0xffff80000009c220
6         init:vfs     0     0xffff800000074000 0xffff80000009c2a8
7         init:fat     0     0xffff80000007a000 0xffff80000009c330
8         tmpfs        0     0xffff800007f36000 0xffff80000009c3b8
9         locfs        0     0xffff800007f76000 0xffff80000009c440
…
26        getterm      0     0xffff800007ab2000 0xffff80000009ccc0
27        bdsh         0     0xffff800007ade000 0xffff80000009cd48
…
38        klog         0     0xffff800000062000 0xffff80000009c110
39        file_bd      0     0xffff800007120000 0xffff800007bc12a8
41        mfs          0     0xffff8000071fc000 0xffff800007bc13b8
```

Analysis

- List all processes

```
kconsole> tasks
[id    ] [name     ] [ctn] [address          ] [as               ]
1         kernel       0     0xffff80000004a000 0xffff80000009c000
2         init:ns      0     0xffff80000005c000 0xffff80000009c088
4         init:loc     0     0xffff800000068000 0xffff80000009c198
5         init:rd      0     0xffff80000006c000 0xffff80000009c220
6         init:vfs     0     0xffff800000074000 0xffff80000009c2a8
7         init:fat     0     0xffff80000007a000 0xffff80000009c330
8         tmpfs        0     0xffff800007f36000 0xffff80000009c3b8
9         locfs        0     0xffff800007f76000 0xffff80000009c440
…
26        getterm      0     0xffff800007ab2000 0xffff80000009ccc0
27        bdsh         0     0xffff800007ade000 0xffff80000009cd48
…
38        klog         0     0xffff800000062000 0xffff80000009c110
39        file_bd      0     0xffff800007120000 0xffff800007bc12a8
41        mfs          0     0xffff8000071fc000 0xffff800007bc13b8
```

Analysis

- Inspect IPC state of bdsh

bdsh 27

```
kconsole> ipc 27
[phone id] [calls] [state
0                0 connected to 2 (init:ns)
1                0 connected to 6 (init:vfs)
2                0 connected to 6 (init:vfs)
3                0 connected to 6 (init:vfs)
4                0 connected to 6 (init:vfs)
5                1 connected to 6 (init:vfs)
6                0 connected to 23 (console)
7                0 connected to 23 (console)
8                0 connected to 4 (init:loc)
[call id         ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
--– incoming answers ---
```

Analysis

- Inspect IPC state of bdsh

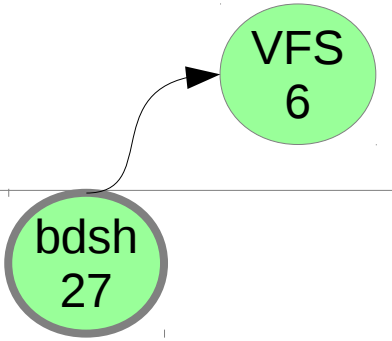


```
kconsole> ipc 27
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 6 (init:vfs)
2                  0 connected to 6 (init:vfs)
3                  0 connected to 6 (init:vfs)
4                  0 connected to 6 (init:vfs)
5                  1 connected to 6 (init:vfs)
6                  0 connected to 23 (console)
7                  0 connected to 23 (console)
8                  0 connected to 4 (init:loc)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
--- incoming answers ---
```

Analysis

- Inspect IPC state of VFS

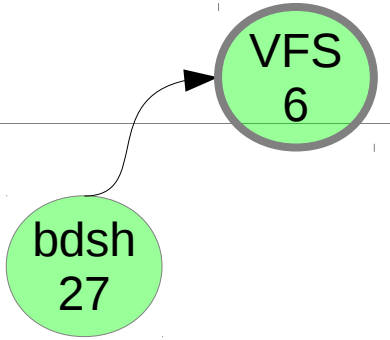


```
kconsole> ipc 6
[phone id] [calls] [state
0                0 connected to 2 (init:ns)
1                0 connected to 7 (init:fat)
2                0 connected to 7 (init:fat)
3                0 connected to 7 (init:fat)
4                0 connected to 8 (tmpfs)
5                0 connected to 9 (locfs)
6                0 connected to 9 (locfs)
7                0 connected to 8 (tmpfs)
8                2 connected to 7 (init:fat)
9                0 connected to 9 (locfs)
10               0 connected to 41 (mfs)
11               0 connected to 41 (mfs)
[call id         ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77800 1032     37     0      0      0      0      0       27 (bdsh)
0xffff800000d77600 1025     3      0      0      0      0      0       39 (file_bd)
--– incoming answers ---
```

Analysis

- Inspect IPC state of VFS

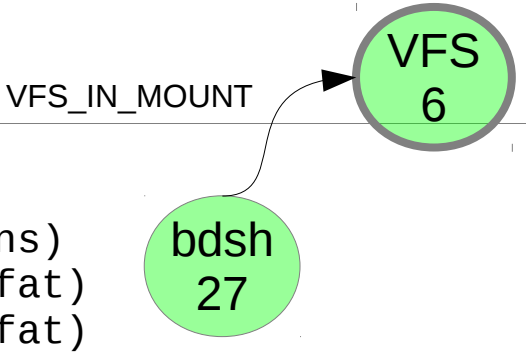


```
kconsole> ipc 6
[phone id] [calls] [state
0                0 connected to 2 (init:ns)
1                0 connected to 7 (init:fat)
2                0 connected to 7 (init:fat)
3                0 connected to 7 (init:fat)
4                0 connected to 8 (tmpfs)
5                0 connected to 9 (locfs)
6                0 connected to 9 (locfs)
7                0 connected to 8 (tmpfs)
8                2 connected to 7 (init:fat)
9                0 connected to 9 (locfs)
10               0 connected to 41 (mfs)
11               0 connected to 41 (mfs)
[call id         ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77800 1032     37     0      0      0      0      0       27 (bdsh)
0xffff800000d77600 1025     3      0      0      0      0      0       39 (file_bd)
--– incoming answers ---
```

Analysis

- Inspect IPC state of VFS

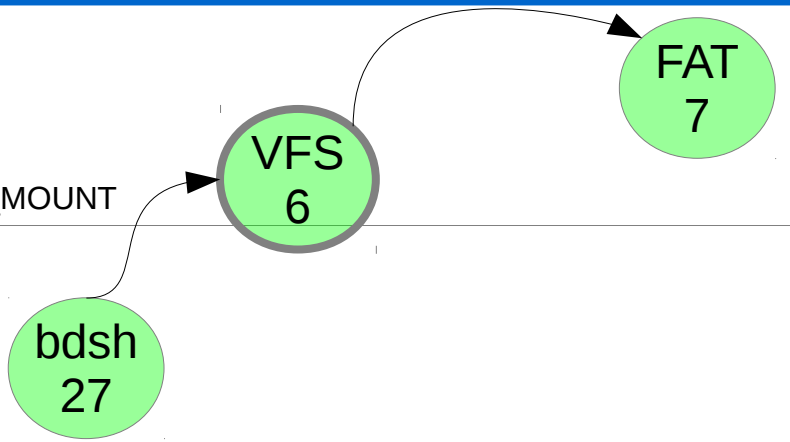


```
kconsole> ipc 6
[phone id] [calls] [state
0                0 connected to 2 (init:ns)
1                0 connected to 7 (init:fat)
2                0 connected to 7 (init:fat)
3                0 connected to 7 (init:fat)
4                0 connected to 8 (tmpfs)
5                0 connected to 9 (locfs)
6                0 connected to 9 (locfs)
7                0 connected to 8 (tmpfs)
8                2 connected to 7 (init:fat)
9                0 connected to 9 (locfs)
10               0 connected to 41 (mfs)
11               0 connected to 41 (mfs)
[call id         ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77800 1032     37     0      0      0      0      0       27 (bdsh)
0xffff800000d77600 1025     3      0      0      0      0      0       39 (file_bd)
--– incoming answers ---
```

Analysis

- Inspect IPC state of FAT

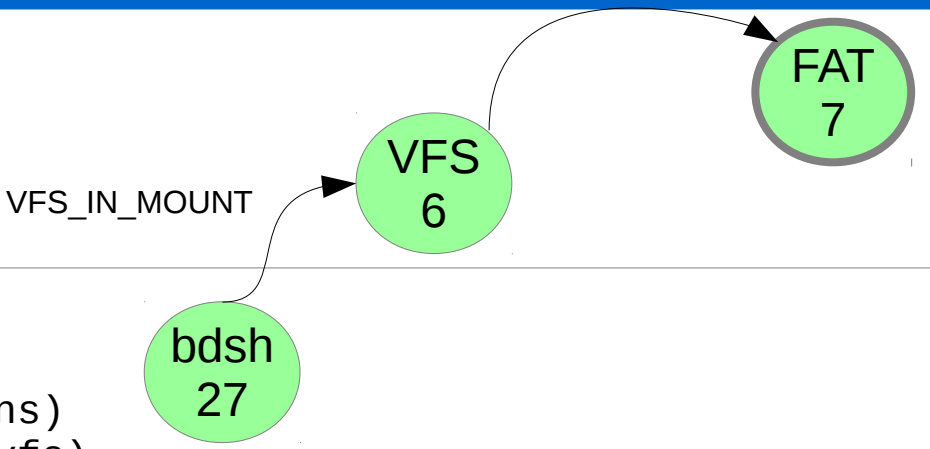


```
kconsole> ipc 7
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 6 (init:vfs)
2                  0 connected to 5 (init:rd)
3                  0 connected to 9 (locfs)
4                  0 connected to 9 (locfs)
5                  0 connected to 8 (tmpfs)
6                  0 connected to 8 (tmpfs)
7                  0 connected to 41 (mfs)
8                  1 connected to 41 (mfs)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77700 1029     10     87     4      37     0      0       6 (init:vfs)
0xffff800000d77500 1025     10     83     16384  0      0      0       6 (init:vfs)
0xffff800000d77900 8        106576 4096   0      0      0      4       39 (file_bd)
--– incoming answers ---
```

Analysis

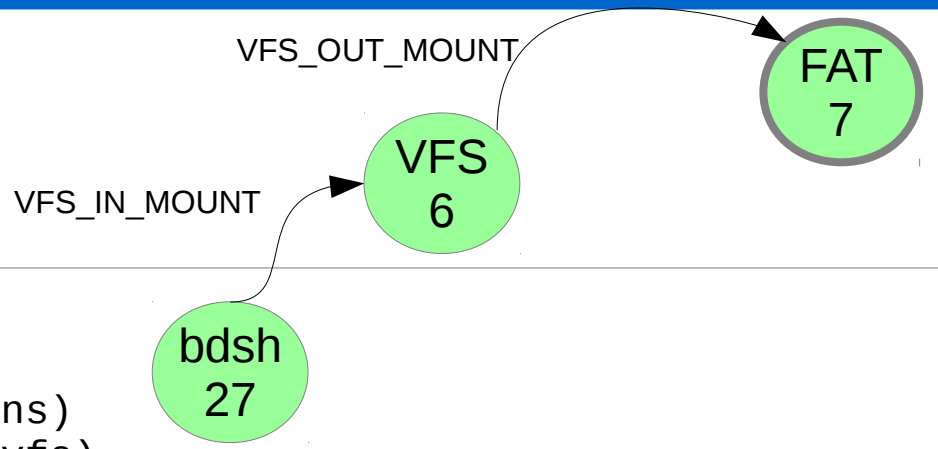

- Inspect IPC state of FAT

```
kconsole> ipc 7
[phone id] [calls] [state
0                0 connected to 2 (init:ns)
1                0 connected to 6 (init:vfs)
2                0 connected to 5 (init:rd)
3                0 connected to 9 (locfs)
4                0 connected to 9 (locfs)
5                0 connected to 8 (tmpfs)
6                0 connected to 8 (tmpfs)
7                0 connected to 41 (mfs)
8                1 connected to 41 (mfs)
[call id         ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77700 1029     10     87     4      37     0      0       6 (init:vfs)
0xffff800000d77500 1025     10     83     16384  0      0      0       6 (init:vfs)
0xffff800000d77900 8        106576 4096   0      0      0      4       39 (file_bd)
--- incoming answers ---
```

Analysis

- Inspect IPC state of FAT

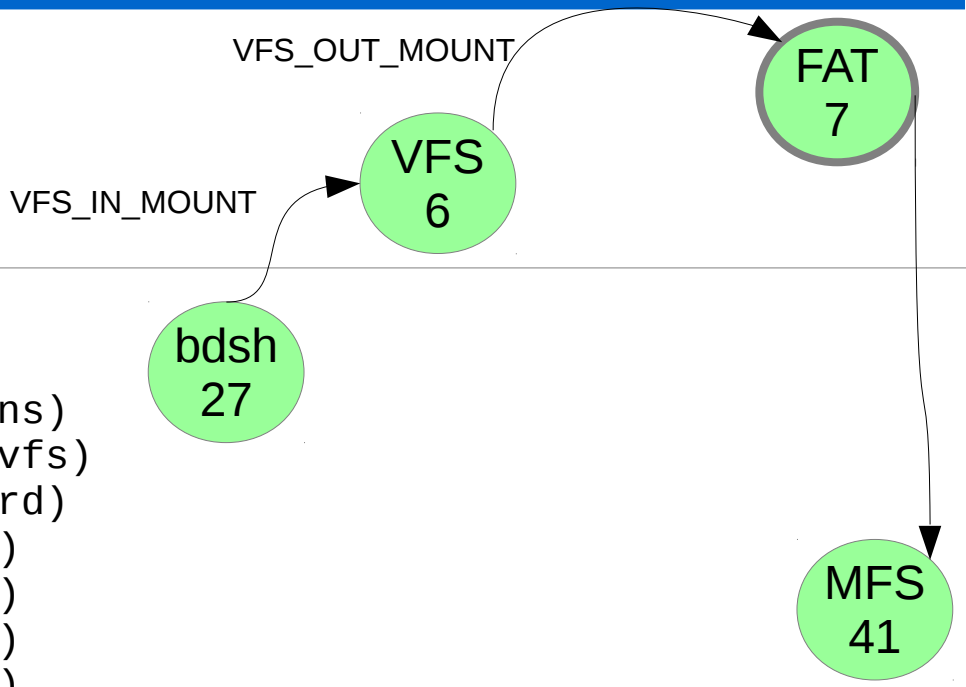


```
kconsole> ipc 7
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 6 (init:vfs)
2                  0 connected to 5 (init:rd)
3                  0 connected to 9 (locfs)
4                  0 connected to 9 (locfs)
5                  0 connected to 8 (tmpfs)
6                  0 connected to 8 (tmpfs)
7                  0 connected to 41 (mfs)
8                  1 connected to 41 (mfs)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77700 1029     10     87     4      37     0      0       6 (init:vfs)
0xffff800000d77500 1025     10     83     16384  0      0      0       6 (init:vfs)
0xffff800000d77900 8        106576 4096   0      0      0      4       39 (file_bd)
--– incoming answers ---
```

Analysis

- Inspect IPC state of MFS

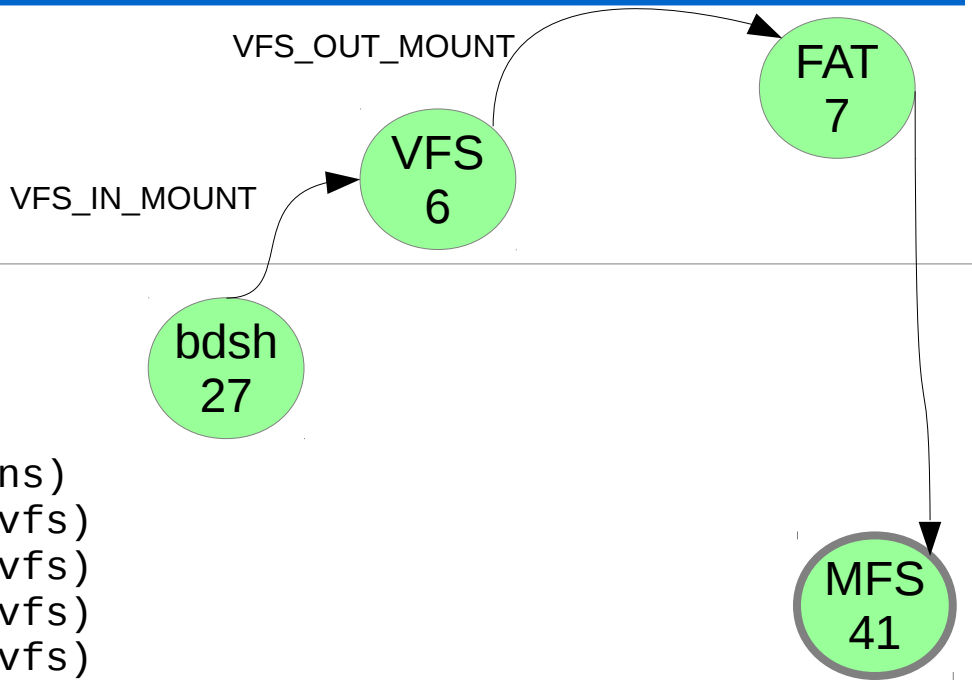


```
kconsole> ipc 41
[phone id] [calls] [state
0                0 connected to 2 (init:ns)
1                0 connected to 6 (init:vfs)
2                0 connected to 6 (init:vfs)
3                0 connected to 6 (init:vfs)
4                0 connected to 6 (init:vfs)
5                0 connected to 6 (init:vfs)
6                0 connected to 6 (init:vfs)
7                1 connected to 39 (file_bd)
[call id         ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77400 1030     37     0      0      0      0      0       7 (init:fat)
--- incoming answers ---
```

Analysis

- Inspect IPC state of MFS

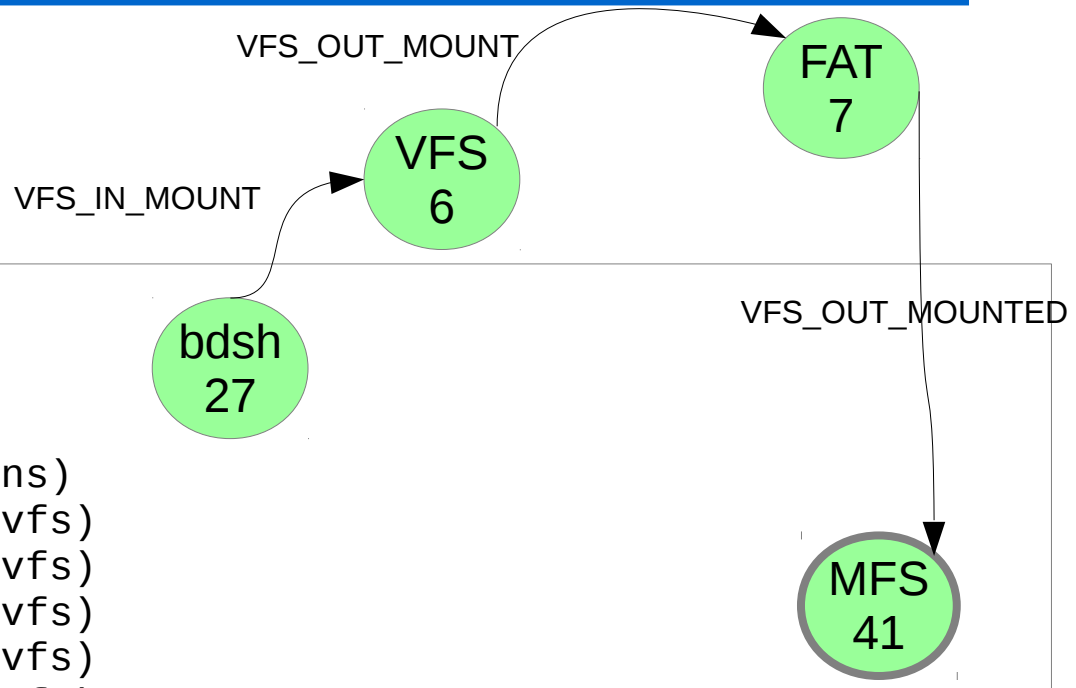


```
kconsole> ipc 41
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 6 (init:vfs)
2                  0 connected to 6 (init:vfs)
3                  0 connected to 6 (init:vfs)
4                  0 connected to 6 (init:vfs)
5                  0 connected to 6 (init:vfs)
6                  0 connected to 6 (init:vfs)
7                  1 connected to 39 (file_bd)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77400 1030     37     0      0      0      0      0       7 (init:fat)
--- incoming answers ---
```

Analysis

- Inspect IPC state of MFS

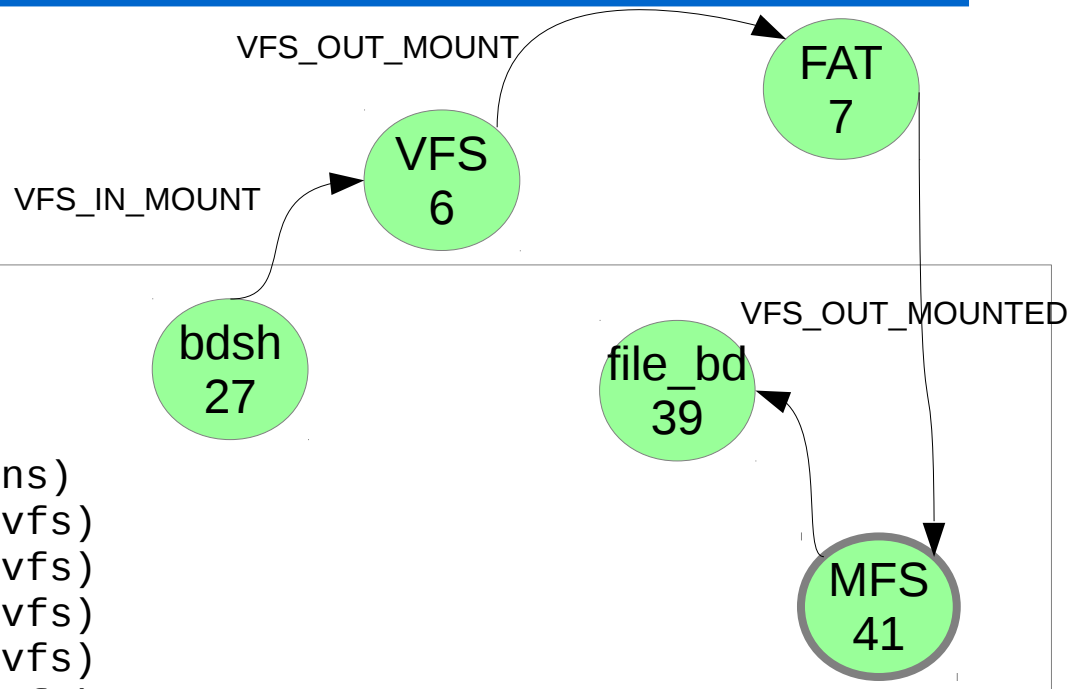


```
kconsole> ipc 41
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 6 (init:vfs)
2                  0 connected to 6 (init:vfs)
3                  0 connected to 6 (init:vfs)
4                  0 connected to 6 (init:vfs)
5                  0 connected to 6 (init:vfs)
6                  0 connected to 6 (init:vfs)
7                  1 connected to 39 (file_bd)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77400 1030     37     0      0      0      0      0       7 (init:fat)
--- incoming answers ---
```

Analysis

- Inspect IPC state of file_bd

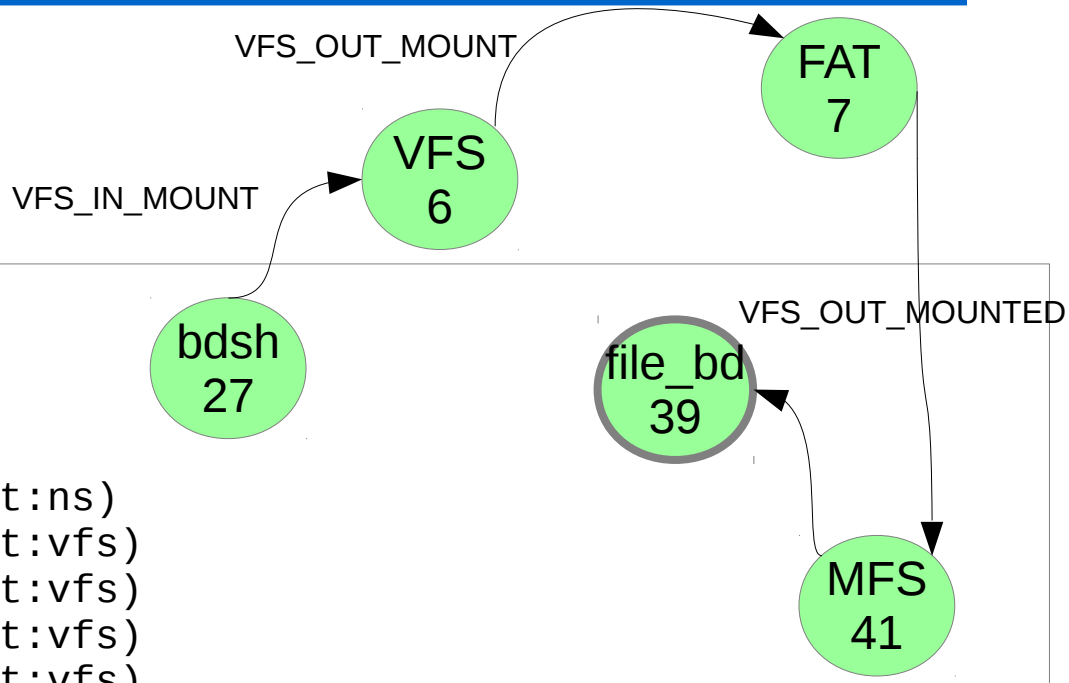


```
kconsole> ipc 39
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 6 (init:vfs)
2                  0 connected to 6 (init:vfs)
3                  0 connected to 6 (init:vfs)
4                  0 connected to 6 (init:vfs)
5                  2 connected to 6 (init:vfs)
6                  0 connected to 4 (init:loc)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77580 1026     32     0      0      0      0      0       41 (mfs)
--- incoming answers ---
```

Analysis

- Inspect IPC state of file_bd

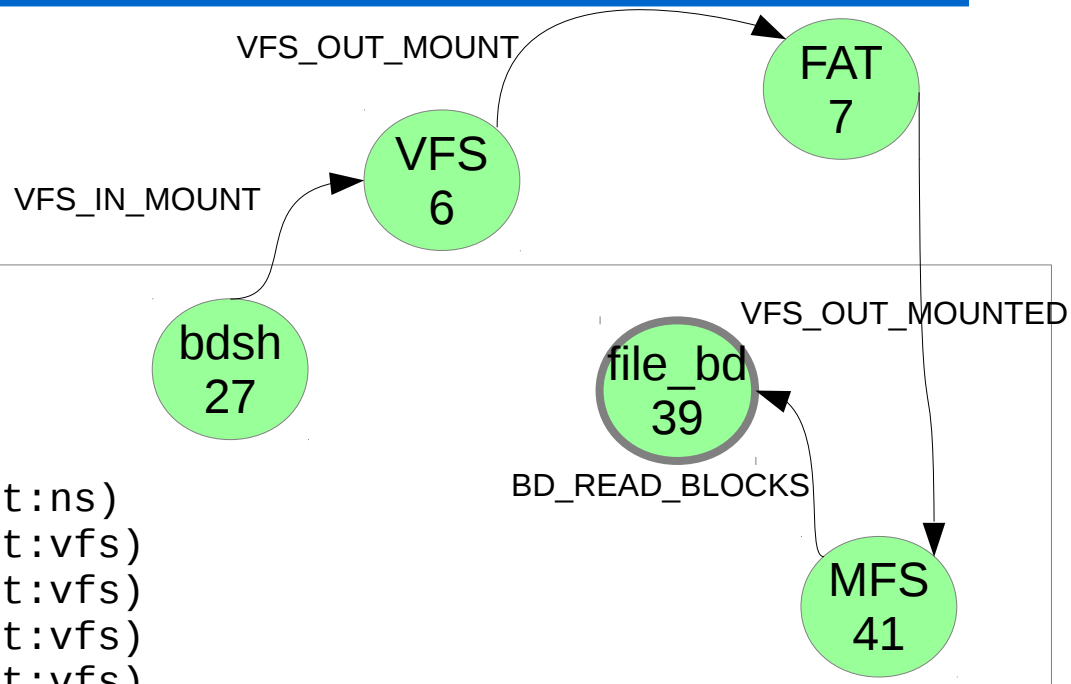


```
kconsole> ipc 39
[phone id] [calls] [state
0                0 connected to 2 (init:ns)
1                0 connected to 6 (init:vfs)
2                0 connected to 6 (init:vfs)
3                0 connected to 6 (init:vfs)
4                0 connected to 6 (init:vfs)
5                2 connected to 6 (init:vfs)
6                0 connected to 4 (init:loc)
[call id         ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77580 1026     32     0      0      0      0      41 (mfs)
--- incoming answers ---
```

Analysis

- Inspect IPC state of file_bd

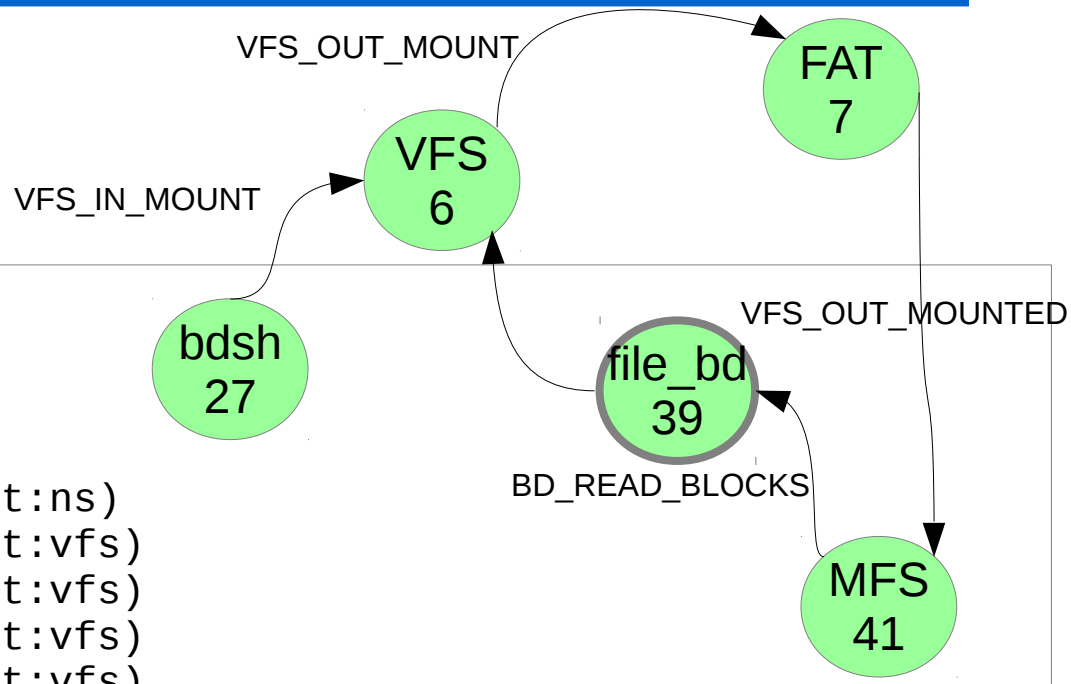


```
kconsole> ipc 39
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 6 (init:vfs)
2                  0 connected to 6 (init:vfs)
3                  0 connected to 6 (init:vfs)
4                  0 connected to 6 (init:vfs)
5                  2 connected to 6 (init:vfs)
6                  0 connected to 4 (init:loc)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77580 1026     32     0      0      0      0      0       41 (mfs)
--- incoming answers ---
```

Analysis

- Re-inspect IPC state of VFS

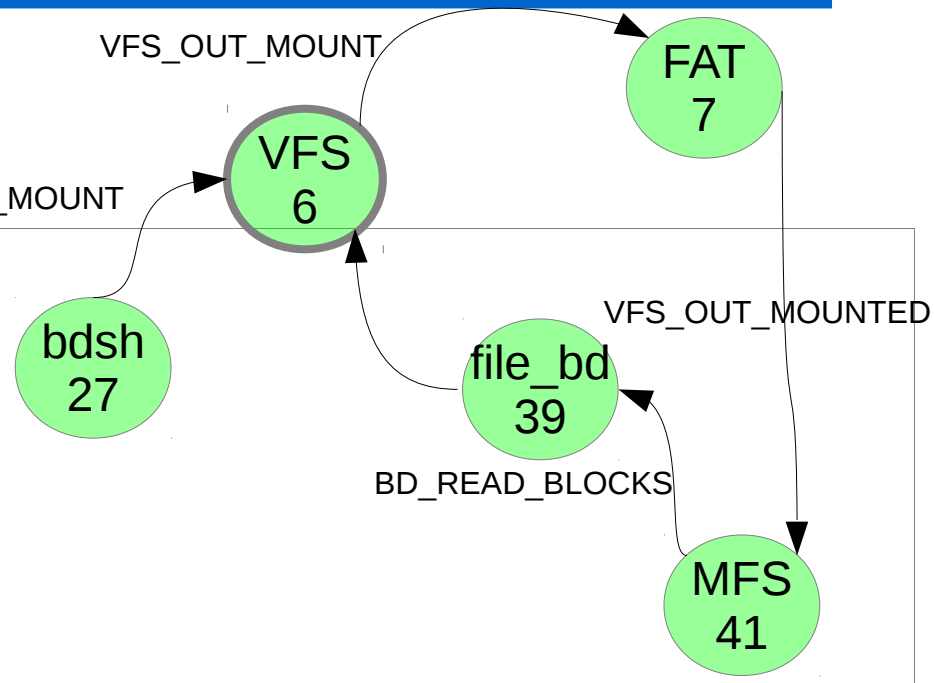


```
kconsole> ipc 6
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 7 (init:fat)
2                  0 connected to 7 (init:fat)
3                  0 connected to 7 (init:fat)
4                  0 connected to 8 (tmpfs)
5                  0 connected to 9 (locfs)
6                  0 connected to 9 (locfs)
7                  0 connected to 8 (tmpfs)
8                  2 connected to 7 (init:fat)
9                  0 connected to 9 (locfs)
10                 0 connected to 41 (mfs)
11                 0 connected to 41 (mfs)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77800 1032     37     0      0      0      0      0       27 (bdsh)
0xffff800000d77600 1025     3      0      0      0      0      0       39 (file_bd)
--– incoming answers ---
```

Analysis

- Re-inspect IPC state of VFS

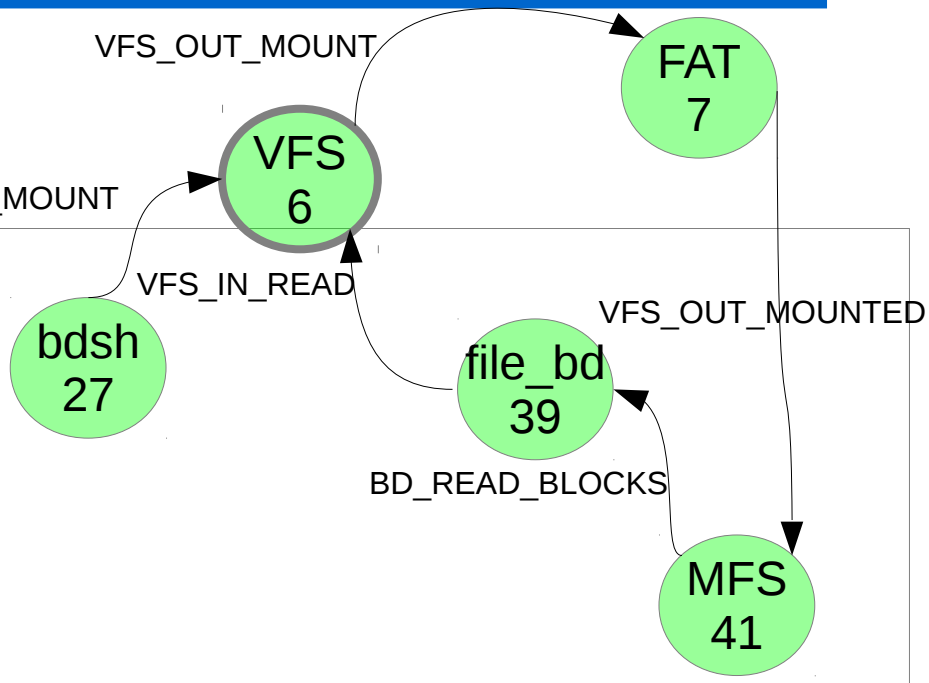


```
kconsole> ipc 6
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 7 (init:fat)
2                  0 connected to 7 (init:fat)
3                  0 connected to 7 (init:fat)
4                  0 connected to 8 (tmpfs)
5                  0 connected to 9 (locfs)
6                  0 connected to 9 (locfs)
7                  0 connected to 8 (tmpfs)
8                  2 connected to 7 (init:fat)
9                  0 connected to 9 (locfs)
10                 0 connected to 41 (mfs)
11                 0 connected to 41 (mfs)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77800 1032     37     0      0      0      0      0       27 (bdsh)
0xffff800000d77600 1025     3      0      0      0      0      0       39 (file_bd)
--– incoming answers ---
```

Analysis

- Re-inspect IPC state of VFS

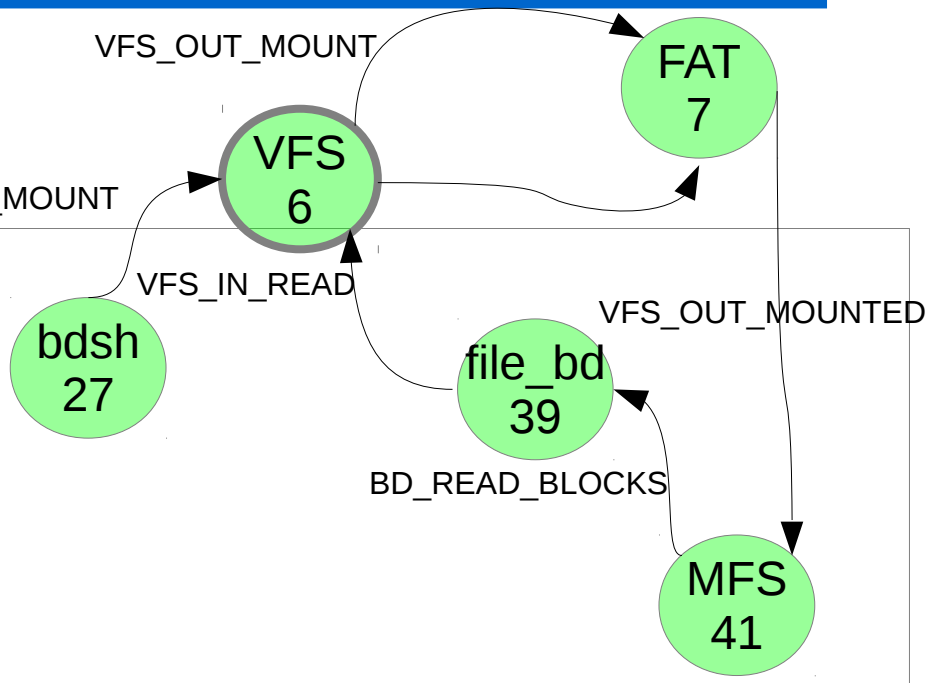


```
kconsole> ipc 6
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 7 (init:fat)
2                  0 connected to 7 (init:fat)
3                  0 connected to 7 (init:fat)
4                  0 connected to 8 (tmpfs)
5                  0 connected to 9 (locfs)
6                  0 connected to 9 (locfs)
7                  0 connected to 8 (tmpfs)
8                  2 connected to 7 (init:fat)
9                  0 connected to 9 (locfs)
10                 0 connected to 41 (mfs)
11                 0 connected to 41 (mfs)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77800 1032     37     0      0      0      0      0       27 (bdsh)
0xffff800000d77600 1025     3      0      0      0      0      0       39 (file_bd)
--– incoming answers ---
```

Analysis

- Re-inspect IPC state of FAT

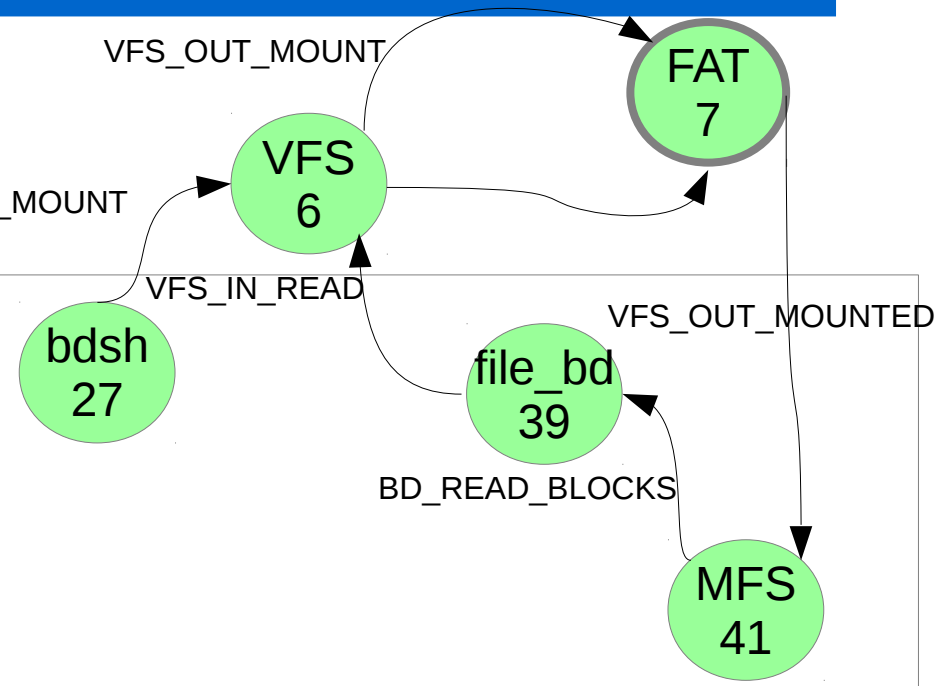


```
kconsole> ipc 7
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 6 (init:vfs)
2                  0 connected to 5 (init:rd)
3                  0 connected to 9 (locfs)
4                  0 connected to 9 (locfs)
5                  0 connected to 8 (tmpfs)
6                  0 connected to 8 (tmpfs)
7                  0 connected to 41 (mfs)
8                  1 connected to 41 (mfs)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77700 1029     10     87     4      37     0      0       6 (init:vfs)
0xffff800000d77500 1025     10     83     16384  0      0      0       6 (init:vfs)
0xffff800000d77900 8        106576 4096   0      0      0      4       39 (file_bd)
--- incoming answers ---
```

Analysis

- Re-inspect IPC state of FAT

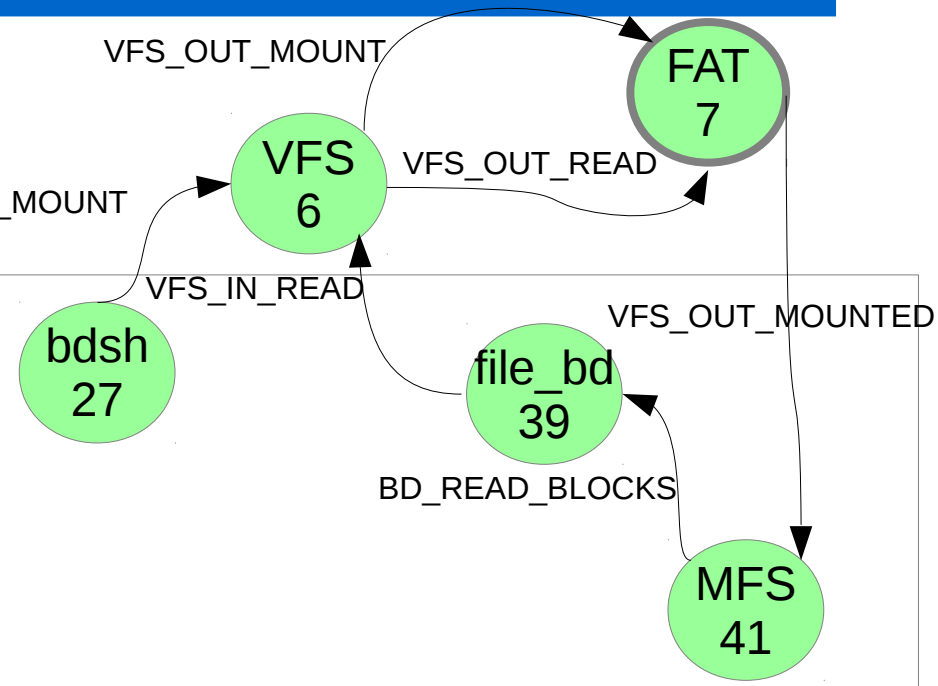


```
kconsole> ipc 7
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 6 (init:vfs)
2                  0 connected to 5 (init:rd)
3                  0 connected to 9 (locfs)
4                  0 connected to 9 (locfs)
5                  0 connected to 8 (tmpfs)
6                  0 connected to 8 (tmpfs)
7                  0 connected to 41 (mfs)
8                  1 connected to 41 (mfs)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77700 1029     10     87     4      37     0      0       6 (init:vfs)
0xffff800000d77500 1025     10     83     16384  0      0      0       6 (init:vfs)
0xffff800000d77900 8        106576 4096   0      0      0      4       39 (file_bd)
--- incoming answers ---
```

Analysis

- Re-inspect IPC state of FAT

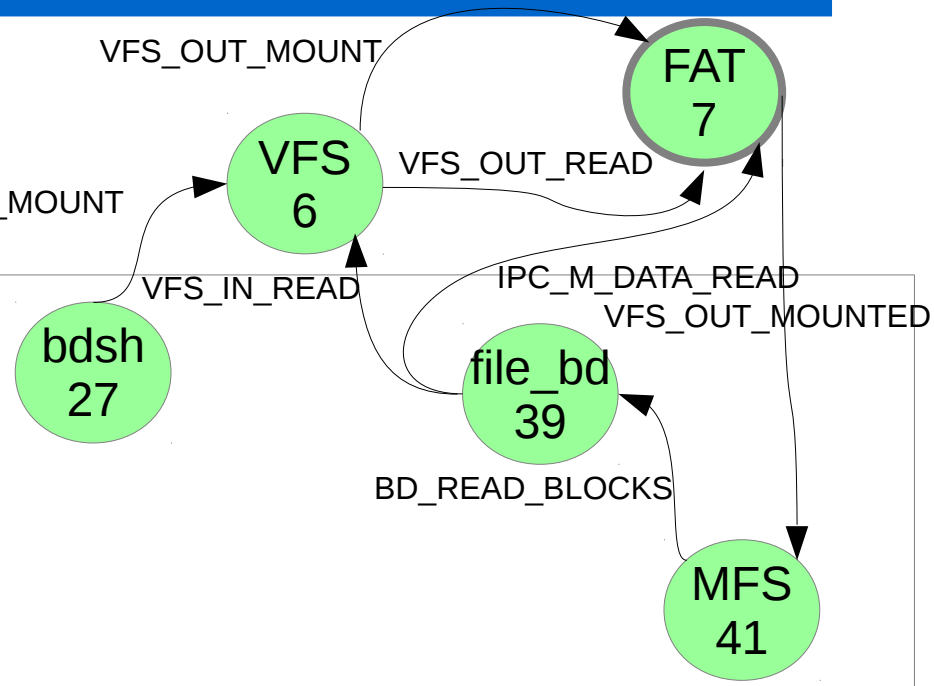


```
kconsole> ipc 7
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 6 (init:vfs)
2                  0 connected to 5 (init:rd)
3                  0 connected to 9 (locfs)
4                  0 connected to 9 (locfs)
5                  0 connected to 8 (tmpfs)
6                  0 connected to 8 (tmpfs)
7                  0 connected to 41 (mfs)
8                  1 connected to 41 (mfs)
[call id           ] [method] [arg1] [arg2] [arg3] [arg4] [arg5] [flags] [sender
--- incoming calls ---
--- dispatched calls ---
0xffff800000d77700 1029     10     87     4      37     0      0       6 (init:vfs)
0xffff800000d77500 1025     10     83     16384  0      0      0       6 (init:vfs)
0xffff800000d77900 8        106576 4096   0      0      0      4       39 (file_bd)
--- incoming answers ---
```

The big picture

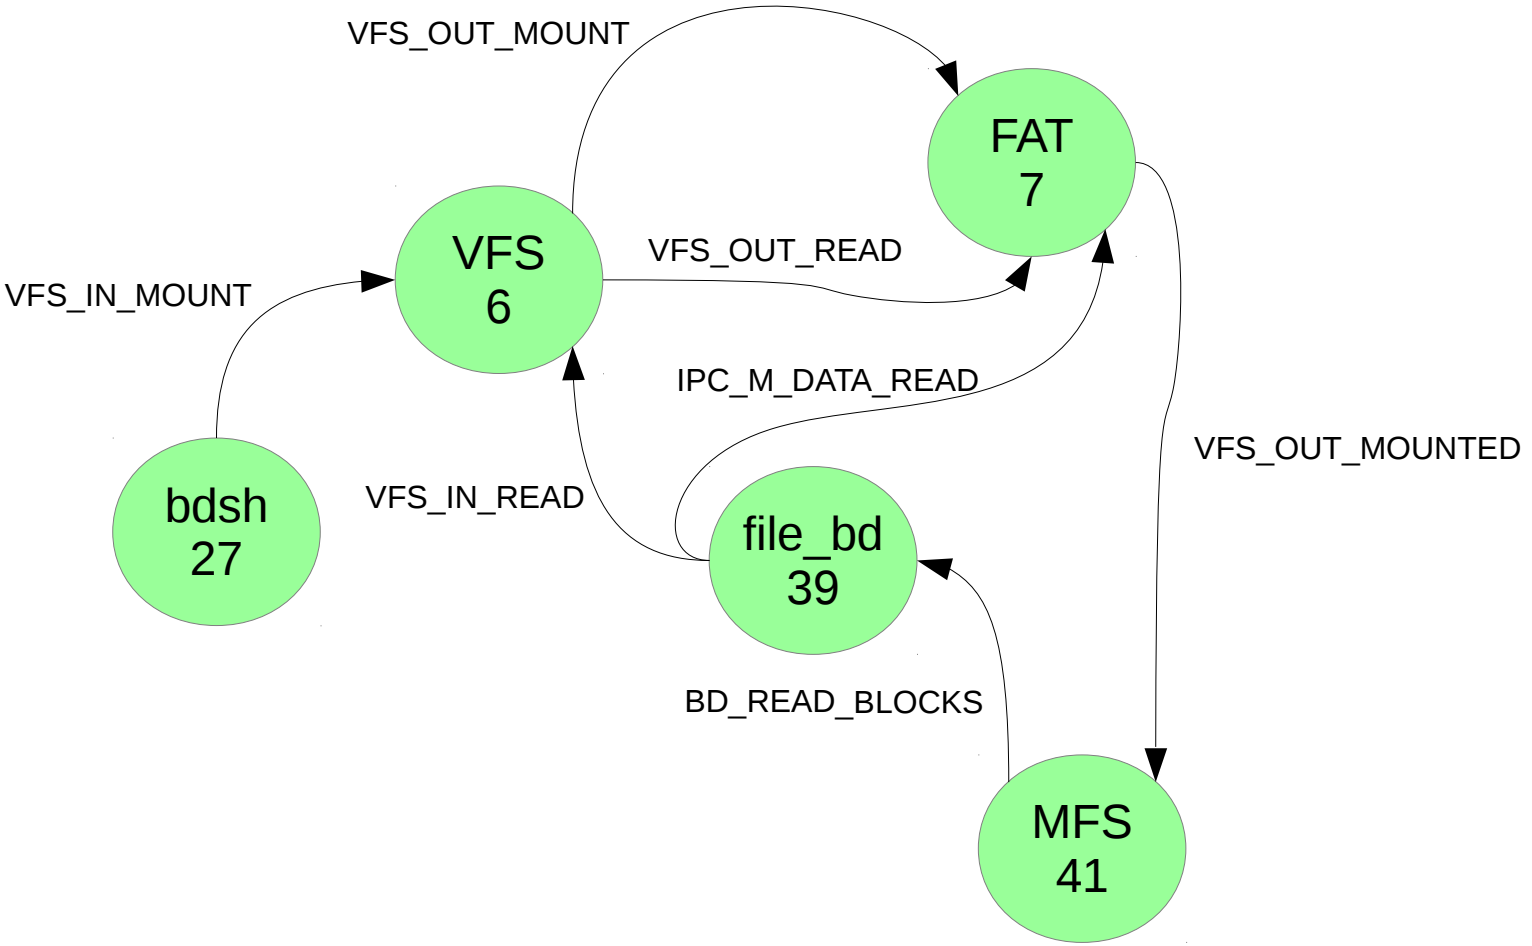

What we DO know

- Everything stopped in FAT
- FAT has 3 unanswered messages
- 2 of those do not correspond to any further IPC
- They just sit there

```
0xffff800000d77700 1029     10     87     4      37     0      0      6 (init:vfs)
0xffff800000d77500 1025     10     83     16384  0      0      0      6 (init:vfs)
0xffff800000d77900 8        106576 4096   0      0      0      4      39 (file_bd)
```

What we DO NOT know

- Why is FAT not processing these calls?

What we NEED to know

- Each IPC connection is handled by a single userspace thread

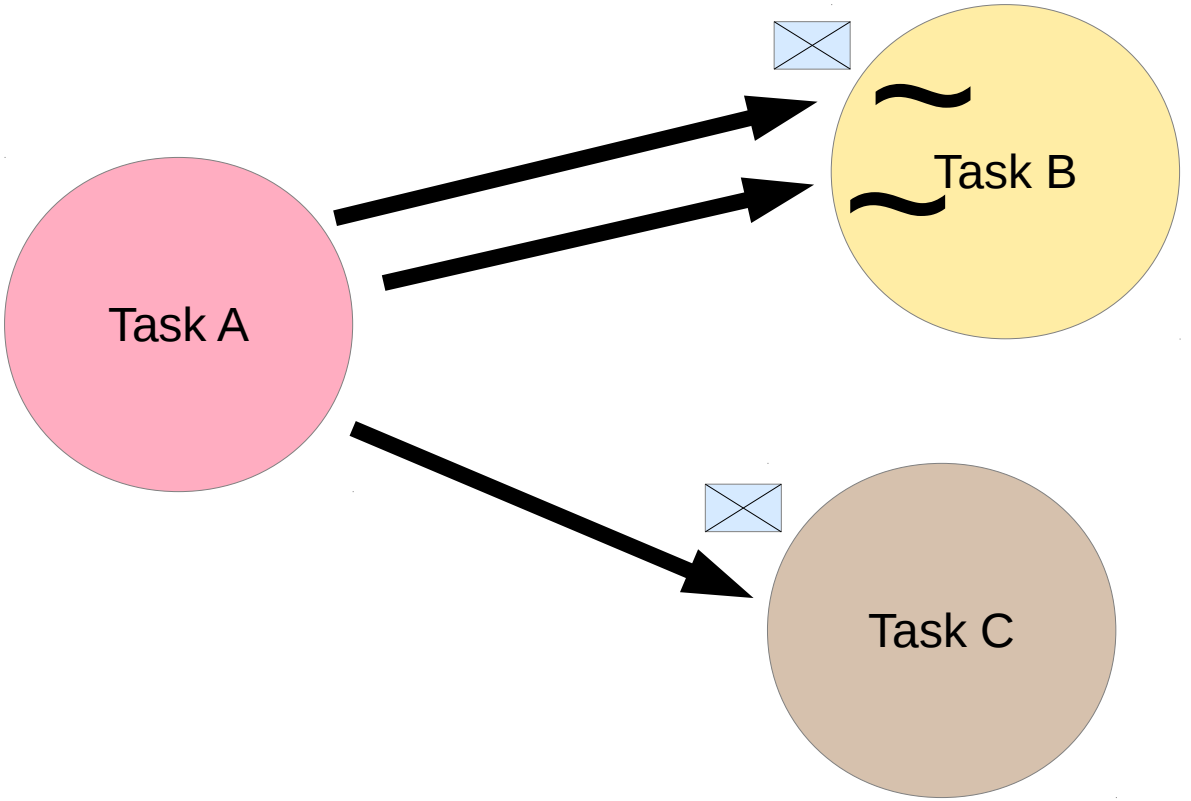

What we need to DO

- Go back to VFS
- Check the IPC connections again

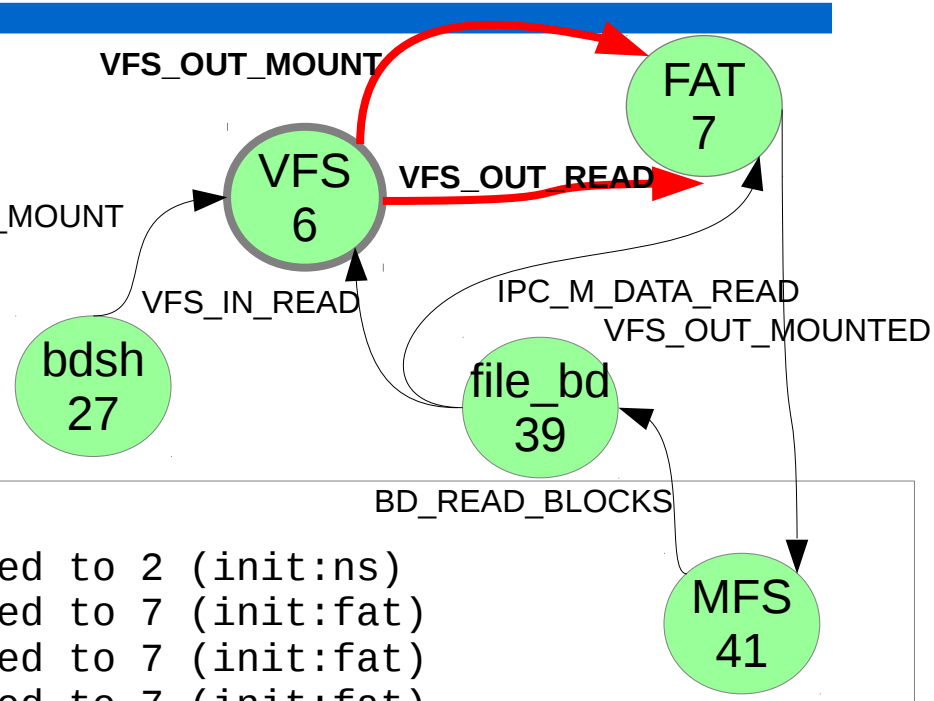


```
[phone id] [calls] [state
0                  0 connected to 2 (init:ns)
1                  0 connected to 7 (init:fat)
2                  0 connected to 7 (init:fat)
3                  0 connected to 7 (init:fat)
4                  0 connected to 8 (tmpfs)
5                  0 connected to 9 (locfs)
6                  0 connected to 9 (locfs)
7                  0 connected to 8 (tmpfs)
8                  2 connected to 7 (init:fat)
9                  0 connected to 9 (locfs)
10                 0 connected to 41 (mfs)
11                 0 connected to 41 (mfs)
```

The reduced picture

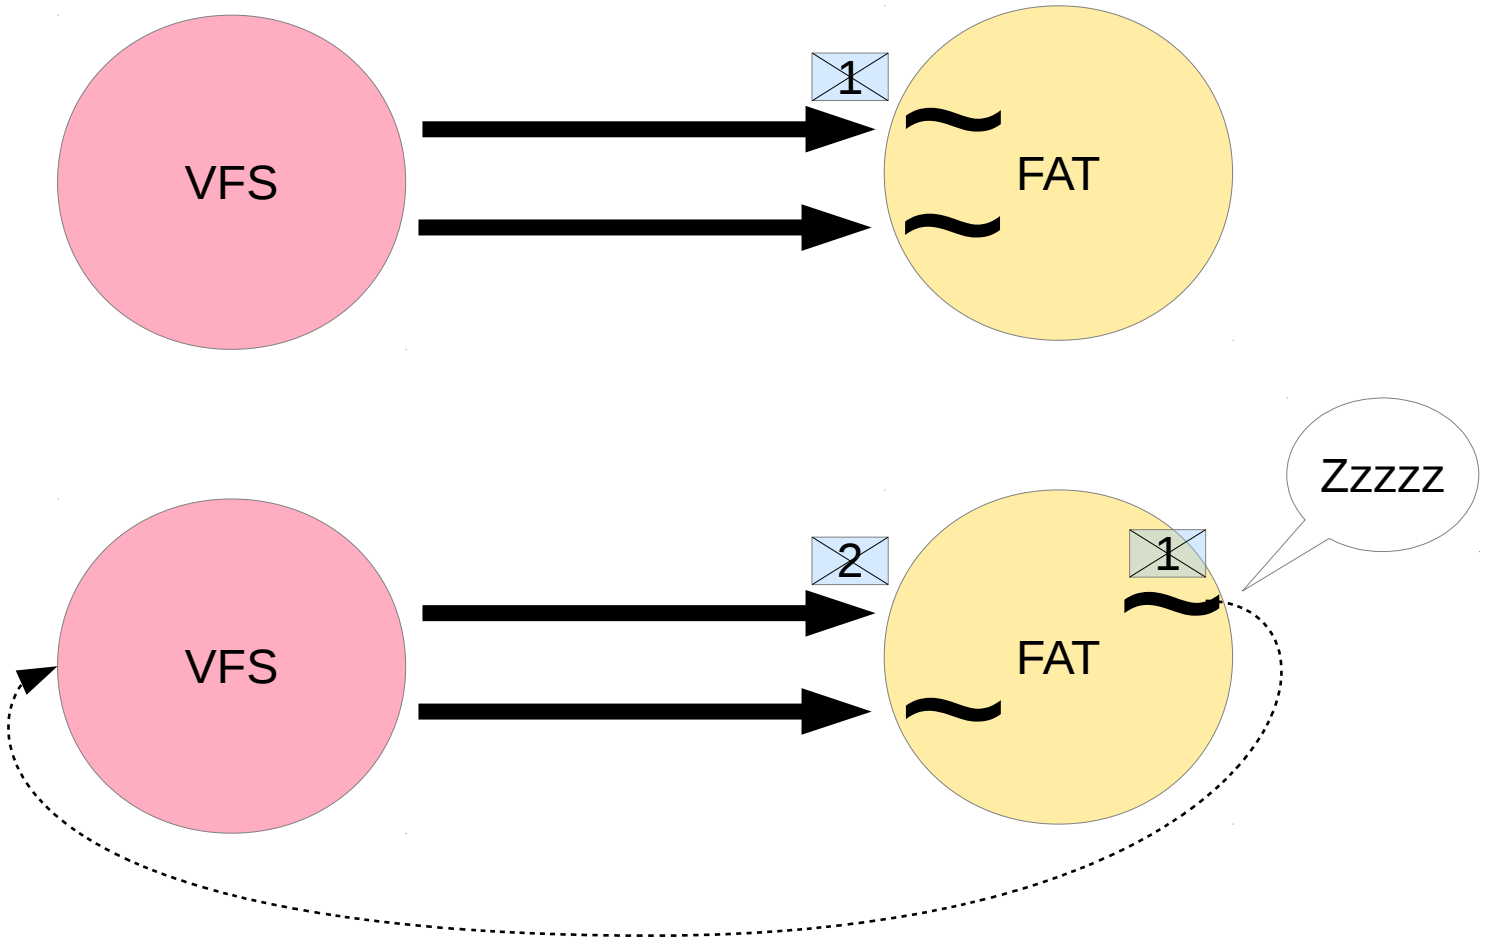

The culprit and the fix

```
revno: 1220
…
timestamp: Fri 2011-09-09 17:50:00 +0200
message:
  Fix deadlock caused by a too early released exchange.
diff:
=== modified file 'uspace/srv/vfs/vfs_ops.c'
--- uspace/srv/vfs/vfs_ops.c    2011-08-19 08:58:50 +0000
+++ uspace/srv/vfs/vfs_ops.c    2011-09-09 15:50:00 +0000
@@ -223,8 +223,14 @@
         return;
     }

+   /*
+    * Wait for the answer before releasing the exchange to avoid deadlock
+    * in case the answer depends on further calls to the same file system.
+    * Think of a case when mounting a FS on a file_bd backed by a file on
+    * the same FS.
+    */
+   async_wait_for(msg, &rc);
    vfs_exchange_release(exch);
-   async_wait_for(msg, &rc);
```

Thank you!

http://www.helenos.org

http://trac.helenos.org/ticket/373

http://jakubsuniversalblog.blogspot.cz/2011/09/debugging-file-system-hang-using.html