

Computer Science CSCI 251

Systems and Networks

Dr. Peter Walsh

Department of Computer Science

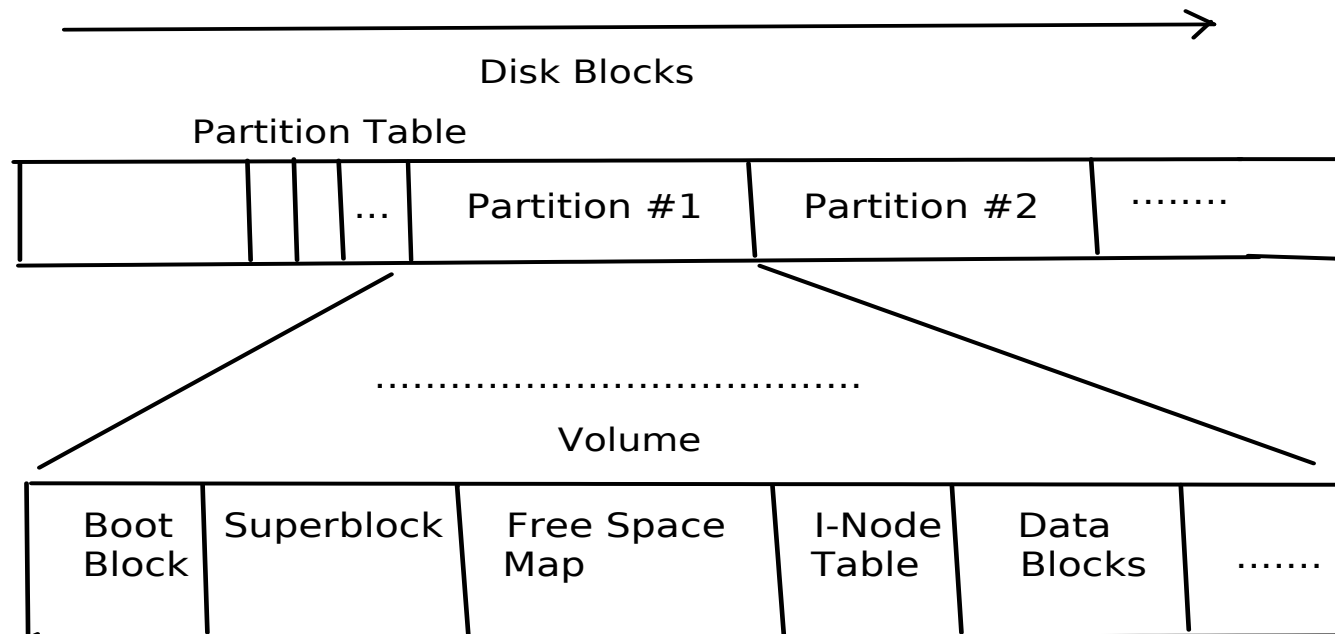
Vancouver Island University

peter.walsh@viu.ca

Disk File Systems

○ Volume

- a partition that contains a file system




Disk File Systems cont.

- Boot Block
 - contains information needed to boot an OS located on the volume (should one exist)
- Superblock
 - balance sheet for the file system meta-data
- Free Space Map
 - a bit-map where each bit is used to indicate whether a particular "object" is free or in-use
- I-Node Table
 - data structure with entries that hold meta-data for a particular file
- Data Blocks
 - hold file and directory data

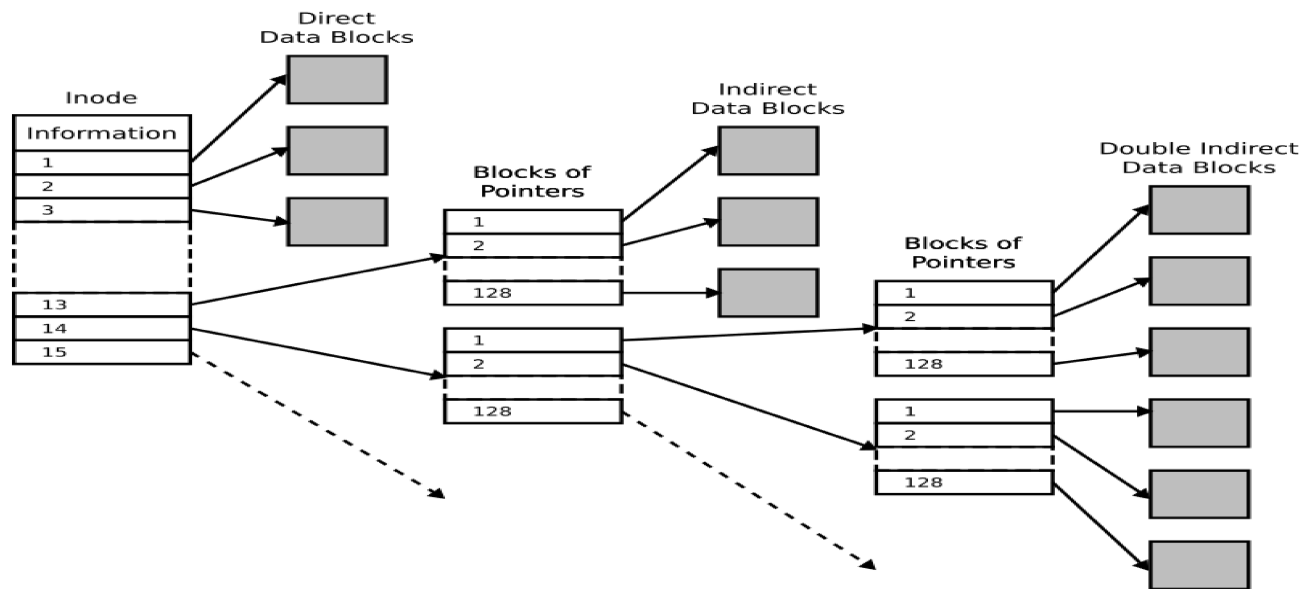
I-Nodes

- I-Node Table
 - array of fixed length I-Nodes
- I-Node Number
 - index into the I-Node Table



I-Node Pointer Structure

○ https://en.wikipedia.org/wiki/Inode_pointer_structure



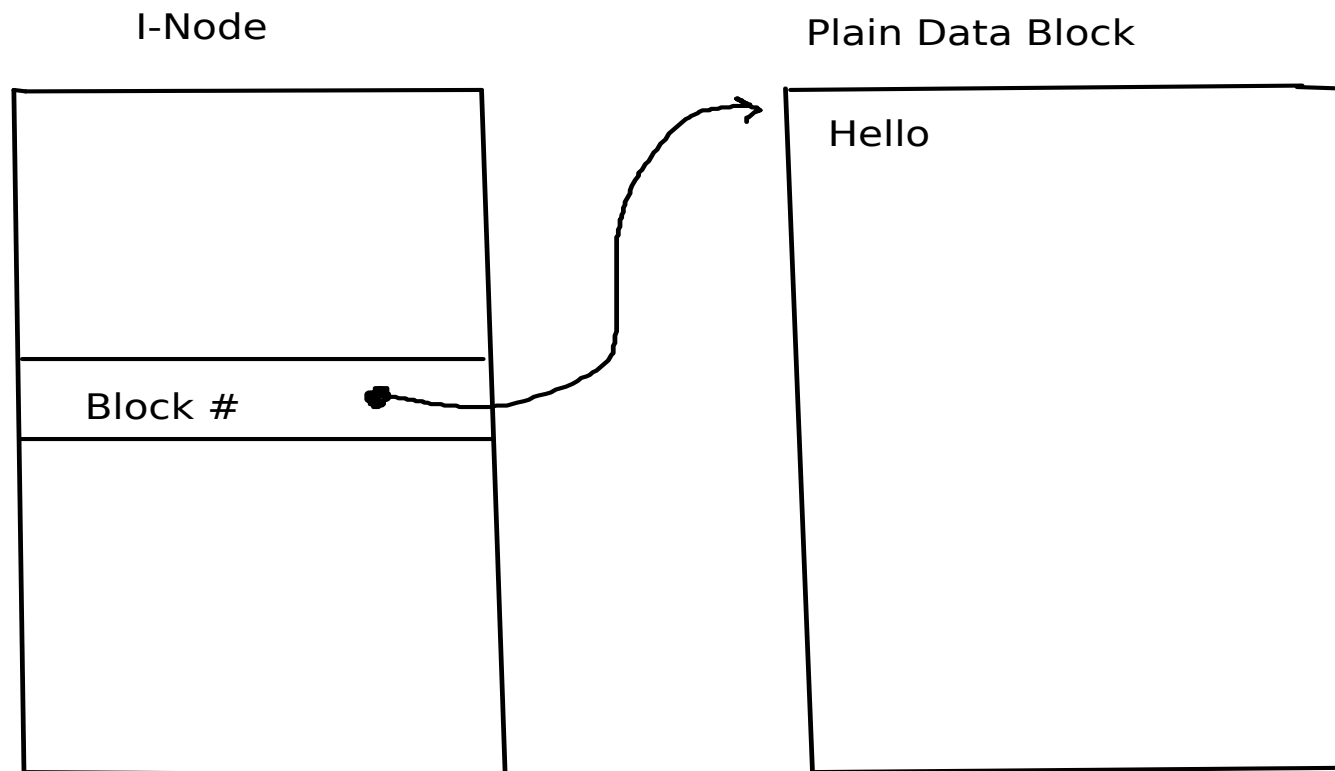
By timtjtim - Own work, CC BY-SA 4.0,
<https://commons.wikimedia.org/w/index.php?curid=75836001>

Data Blocks

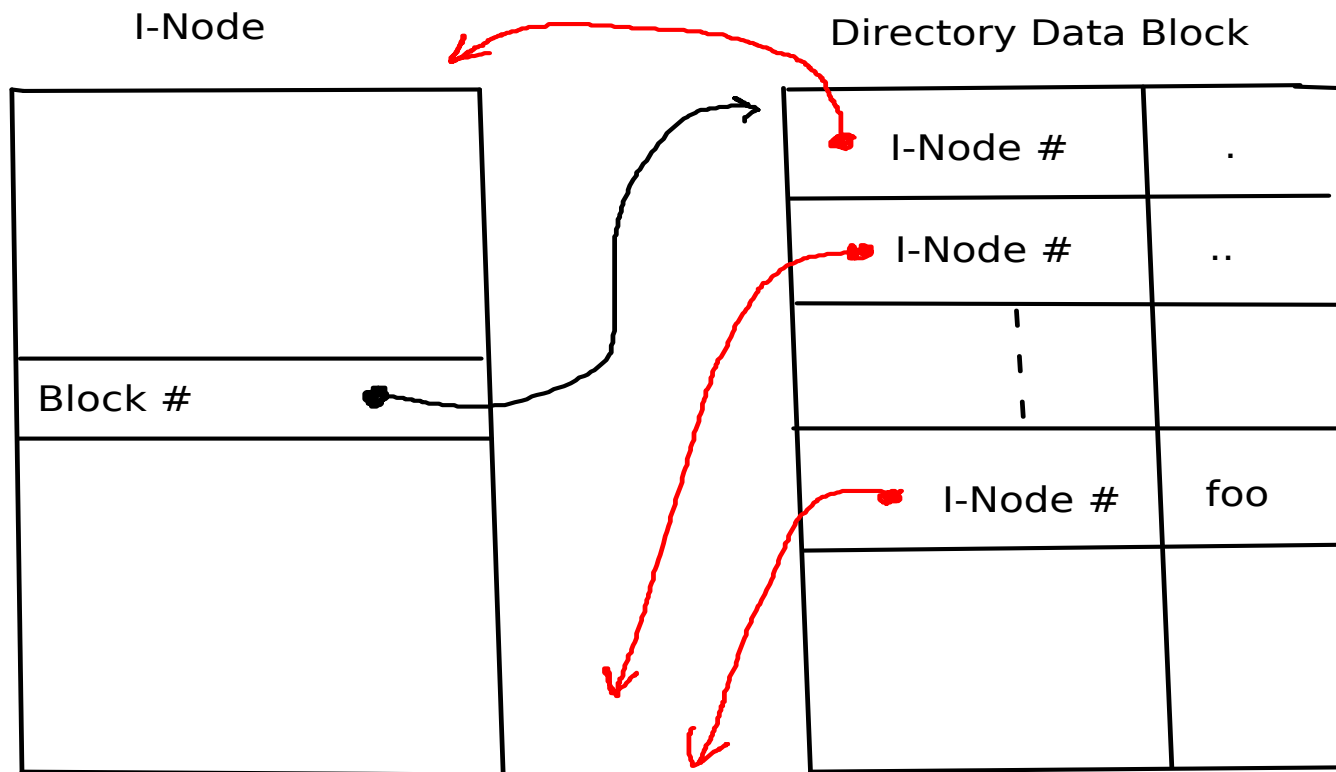
An I-Node can reference three types of data blocks.

- Data Block Types
 - Plain Data
 - Directory Data
 - Symbolic Link Data

Plain Data Blocks



Directory Data Blocks




Access /home/peter/HelloWorld

- the kernel knows the I-Node number of the root directory and the current directory (`pwd`)
- the kernel fetches from I-Node 2, the pointer(s) to the Data Block(s) for the root directory
- the kernel searches the Data Block(s) for `home` and extracts its I-Node number n
- the kernel fetches from I-Node n , the pointer(s) to the Data Block(s) for the `/home` directory
- the kernel searches the Data Block(s) for `peter` and extracts its I-Node number m
- the kernel fetches from I-Node m , the pointer(s) to the Data Block(s) for the `/home/peter` directory
- the kernel searches the Data Block(s) for `HelloWorld` and extracts its I-Node number o
- the kernel fetches from I-Node o , the pointer(s) to the Data Block(s) for the file `/home/peter/HelloWorld`

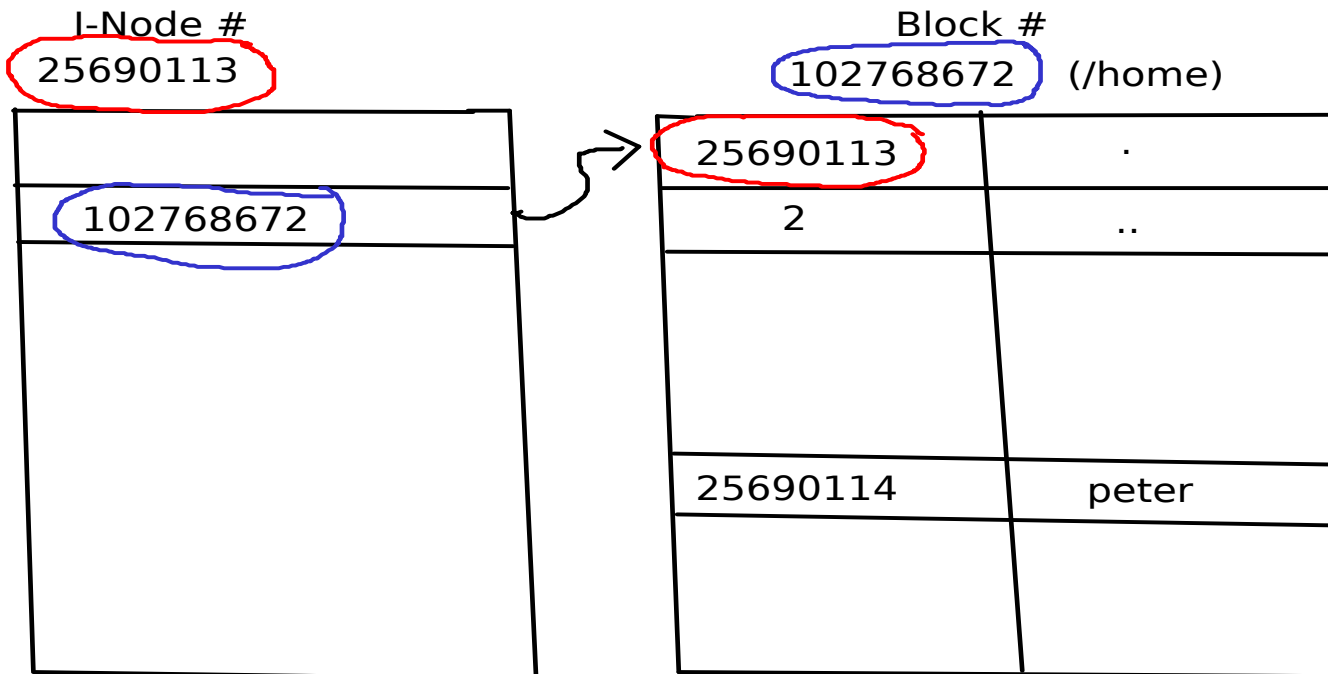
Access /home/peter/HelloWorld

○ Root I-Node




Access /home/peter/HelloWorld cont.

○ /home I-Node



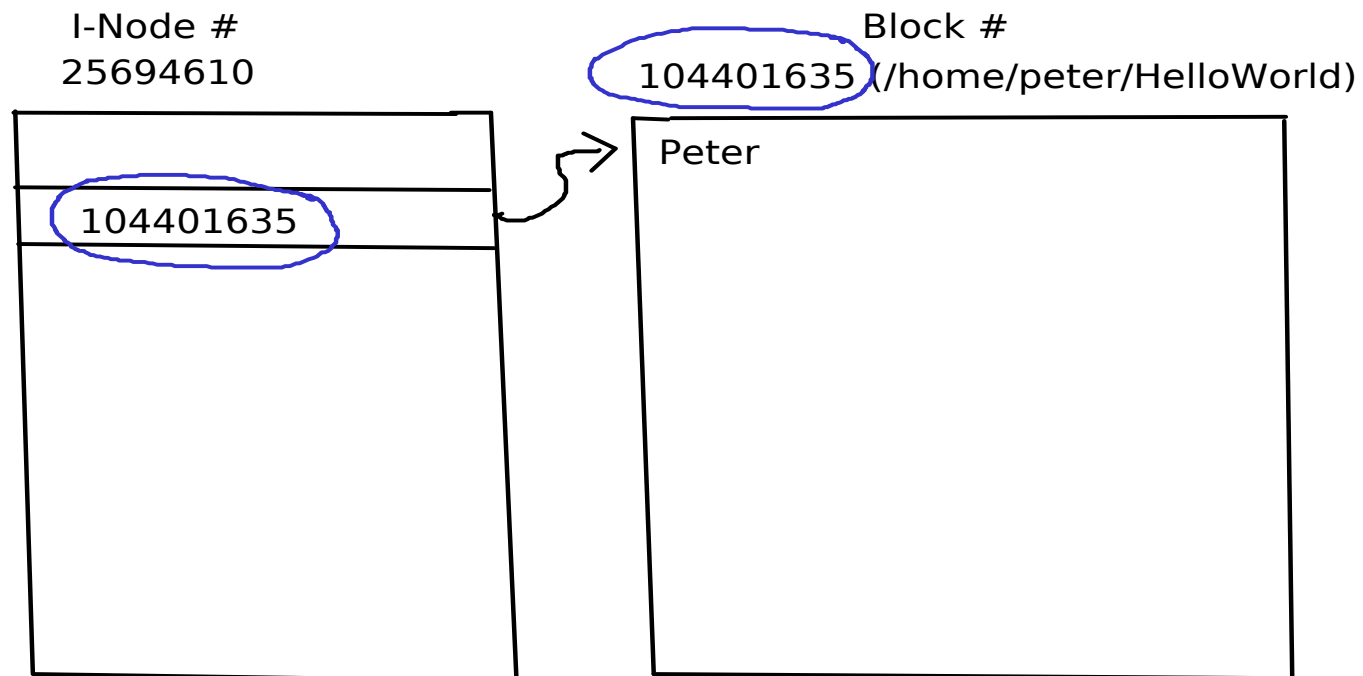
Access /home/peter/HelloWorld cont.

○ /home/peter I-Node



Access /home/peter/HelloWorld cont.

- /home/peter/HelloWorld I-Node



Debug File System

○ Enter Debug Shell (debugfs /dev/sda5)

● debugfs: stat /

Inode: 2 Type: directory Mode: 0755 Flags: 0x80000
Generation: 0 Version: 0x00000000:0000001f
User: 0 Group: 0 Project: 0 Size: 4096
File ACL: 0
Links: 20 Blockcount: 8
Fragment: Address: 0 Number: 0 Size: 0
ctime: 0x5f1f48fe:d3415644 -- Mon Jul 27 14:37:02 2020
atime: 0x60244760:00ecab58 -- Wed Feb 10 12:51:44 2021
mtime: 0x5f1f48fe:d3415644 -- Mon Jul 27 14:37:02 2020
crtime: 0x5edef5b9:00000000 -- Mon Jun 8 19:36:41 2020
Size of extra inode fields: 32
Inode checksum: 0x3ad98372
EXTENTS:
(0):9308

Debug File System cont.

- debugfs: stat /home

```
Inode: 25690113 Type: directory Mode: 0755 Flags: 0x80000  
Generation: 1359808648 Version: 0x00000000:00000002  
User: 0 Group: 0 Project: 0 Size: 4096  
File ACL: 0  
Links: 3 Blockcount: 8  
Fragment: Address: 0 Number: 0 Size: 0  
ctime: 0x5edef64a:8c5b5e40 -- Mon Jun 8 19:39:06 2020  
atime: 0x602447c4:bae35a9c -- Wed Feb 10 12:53:24 2021  
mtime: 0x5edef64a:8c5b5e40 -- Mon Jun 8 19:39:06 2020  
crttime: 0x5edef5d0:96a74864 -- Mon Jun 8 19:37:04 2020  
Size of extra inode fields: 32  
Inode checksum: 0x34021fa2  
EXTENTS:  
(0):102768672
```

Debug File System cont.

- debugfs: stat /home/peter

Inode: 25690114 Type: directory Mode: 0755 Flags: 0x80000
Generation: 2569425769 Version: 0x00000000:00001a82
User: 1000 Group: 1000 Project: 0 Size: 4096
File ACL: 0
Links: 59 Blockcount: 8
Fragment: Address: 0 Number: 0 Size: 0
ctime: 0x602479a1:21147d84 -- Wed Feb 10 16:26:09 2021
atime: 0x602479a8:28b3950c -- Wed Feb 10 16:26:16 2021
mtime: 0x602479a1:21147d84 -- Wed Feb 10 16:26:09 2021
crttime: 0x5edef64a:8c5b5e40 -- Mon Jun 8 19:39:06 2020
Size of extra inode fields: 32
Inode checksum: 0x9842391a
EXTENTS:
(0) 102768673

Debug File System cont.

- debugfs: stat /home/peter/HelloWorld

```
Inode: 25694610 Type: regular Mode: 0664 Flags: 0x80000
Generation: 3814972822 Version: 0x00000000:00000001
User: 1000 Group: 1000 Project: 0 Size: 6
File ACL: 0
Links: 1 Blockcount: 8
Fragment: Address: 0 Number: 0 Size: 0
ctime: 0x6022f1d4:60c983dc -- Tue Feb 9 12:34:28 2021
atime: 0x60232406:69b02224 -- Tue Feb 9 16:08:38 2021
mtime: 0x6022f1d4:60c983dc -- Tue Feb 9 12:34:28 2021
crttime: 0x6022f1d4:60c983dc -- Tue Feb 9 12:34:28 2021
Size of extra inode fields: 32
Inode checksum: 0x9f15f3c7
EXTENTS:
(0):104401635
```

File Contents

```
peter@cobra:~$ cat /home/peter/HelloWorld
```

```
Peter
```

```
debugfs: block_dump 104401635  
0000 5065 7465 720a 0000 0000 0000 0000 0000 0000 Peter.....
```

```
root@cobra:~# dd if=/dev/sda5 bs=4096 skip=104401635
```

```
count=1 | hexdump -C
```

```
1+0 records in
```

```
1+0 records out
```

```
00000000 50 65 74 65 72 0a 00 00 00 00 00 00 00 00 00 |Peter...|
```

```
4096 bytes (4.1 kB, 4.0 KiB) copied, 9.421e-05 s, 43.5 MB/s
```

```
root@cobra:/home/peter# dd if=/dev/sda5 bs=512 skip=835213080
```

```
count=1 | hexdump -C
```

```
1+0 records in
```

```
1+0 records out
```

```
512 bytes copied, 0.000122351 s, 4.2 MB/s
```

```
00000000 50 65 74 65 72 0a 00 00 00 00 00 00 00 00 00 |Peter..|
```

Hard File Links

- two files that share the same I-Node

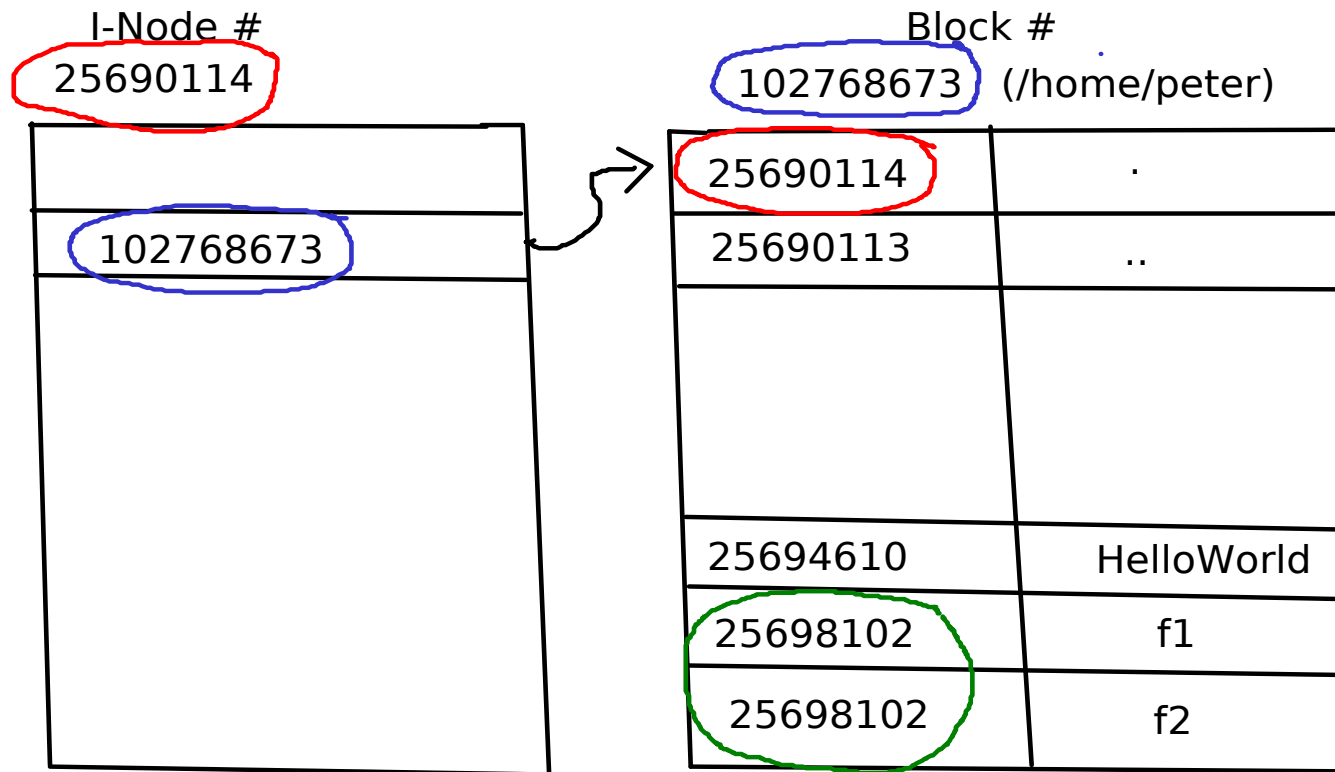
```
# Assume ext4 file system on /dev/sda5  
# (Root File System)
```

```
peter@cobra:~$ echo "Peter Again" > f1  
peter@cobra:~$ ln f1 f2
```

- modifications to `f1` are reflected in `f2` and vice versa
- remove `f1`, `f2` can still access the data
- files must be in the same file system


Access /home/peter/f1 And f2

○ /home/peter I-Node



Access /home/peter/f1 And f2 cont.

- /home/peter/f1 and f2 I-Node



Debug File System

- Enter Debug Shell (debugfs /dev/sda5)
 - debugfs: stat /home/peter/f1

```
Inode: 25698102 Type: regular Mode: 0664 Flags: 0x80000
Generation: 842322228 Version: 0x00000000:00000001
User: 1000 Group: 1000 Project: 0 Size: 12
File ACL: 0
Links: 2 Blockcount: 8
Fragment: Address: 0 Number: 0 Size: 0
ctime: 0x6025533c:e006a510 -- Thu Feb 11 07:54:36 2021
atime: 0x60255336:65006e00 -- Thu Feb 11 07:54:30 2021
mtime: 0x60255336:65006e00 -- Thu Feb 11 07:54:30 2021
crtime: 0x60255336:65006e00 -- Thu Feb 11 07:54:30 2021
Size of extra inode fields: 32
Inode checksum: 0xe3650674
EXTENTS:
(0):102800685
```

Debug File System cont.

- debugfs: `stat /home/peter/f2`

```
Inode: 25698102 Type: regular Mode: 0664 Flags: 0x80000
Generation: 842322228 Version: 0x00000000:00000001
User: 1000 Group: 1000 Project: 0 Size: 12
File ACL: 0
Links: 2 Blockcount: 8
Fragment: Address: 0 Number: 0 Size: 0
 ctime: 0x6025533c:e006a510 -- Thu Feb 11 07:54:36 2021
 atime: 0x60255336:65006e00 -- Thu Feb 11 07:54:30 2021
 mtime: 0x60255336:65006e00 -- Thu Feb 11 07:54:30 2021
 crtime: 0x60255336:65006e00 -- Thu Feb 11 07:54:30 2021
Size of extra inode fields: 32
Inode checksum: 0xe3650674
EXTENTS:
(0):102800685
```

File Contents

```
peter@cobra:~$ cat f1  
Peter Again
```

```
peter@cobra:~$ cat f2  
Peter Again
```

```
debugfs: block_dump 102800685  
0000 5065 7465 7220 4167 6169 6e0a 0000 0000 Peter Again.....  
0020 0000 0000 0000 0000 0000 0000 0000 0000 .....
```

```
root@cobra:/home/peter# dd if=/dev/sda5 bs=4096  
skip=102800685 count=1 | hexdump -C  
1+0 records in  
1+0 records out  
00000000 50 65 74 65 72 20 41 67 61 69 6e 0a 00 00 00 00 Peter Again..
```


Symbolic (Soft) File Links

- a file that contains a pointer to another file

```
# Assume ext4 file system on /dev/sda5 (Root File System)
```

```
# Assume ext3 file system on /dev/sdb2 (Thumb Drive)
```

```
peter@cobra:~$ echo "Peter Again Again" > g1
```

```
root@cobra:/media/peter/ed06746b-..# ln -s /home/peter/g1 g2
```

- modifications to g1 are reflected in g2 and vice versa
- remove g1, g2 can NOT access the data
- files may be on different file systems

Debug File System

- Enter Debug Shell (debugfs /dev/sdb2)
 - debugfs: stat g2

```
Inode: 14 Type: symlink Mode: 0777 Flags: 0x0
Generation: 2962796488 Version: 0x00000000:00000001
User: 0 Group: 0 Project: 0 Size: 14
File ACL: 0
Links: 1 Blockcount: 0
Fragment: Address: 0 Number: 0 Size: 0
ctime: 0x602570c3:b92a4d80 -- Thu Feb 11 10:00:35 2021
atime: 0x602570cf:c68f8f40 -- Thu Feb 11 10:00:47 2021
mtime: 0x602570c3:b92a4d80 -- Thu Feb 11 10:00:35 2021
crtime: 0x602570c3:b92a4d80 -- Thu Feb 11 10:00:35 2021
Size of extra inode fields: 32
Fast link dest: "/home/peter/g1"
(END)
```

Debug File System cont.

- debugfs: inode_dump <14>

```
0000 ffa1 0000 0e00 0000 cf70 2560 c370 2560 .....p%`.p%`
0020 c370 2560 0000 0000 0000 0100 0000 0000 .p%`.....
0040 0000 0000 0100 0000 2f68 6f6d 652f 7065 ...../home/pe
0060 7465 722f 6731 0000 0000 0000 0000 0000 ter/g1.....
0100 0000 0000 0000 0000 0000 0000 0000 0000 .....
*
```