Edinburgh Napier UNIVERSITY

CSN08101
Digital Forensics
Lecture 4: System Level Disk Control

Module Leader: Dr Gordon Russell
Lecturers: Robert Ludwiniak

Edinburgh Napier UNIVERSITY

Essential Linux for Forensics

You will learn in this lecture:

- Command Summary:
  - sudo
  - mount/umount
  - fdisk family (cfdisk)
  - dd
  - strings
  - dmsetup
  - losetup
  - blockdev
- This week is all about:
  - Doing things as the administrator
  - Reading and Writing disk blocks
  - Accessing partition information
  - Mounting and unmounting and loop devices
  - Snapshots

Edinburgh Napier UNIVERSITY

Directory Tree

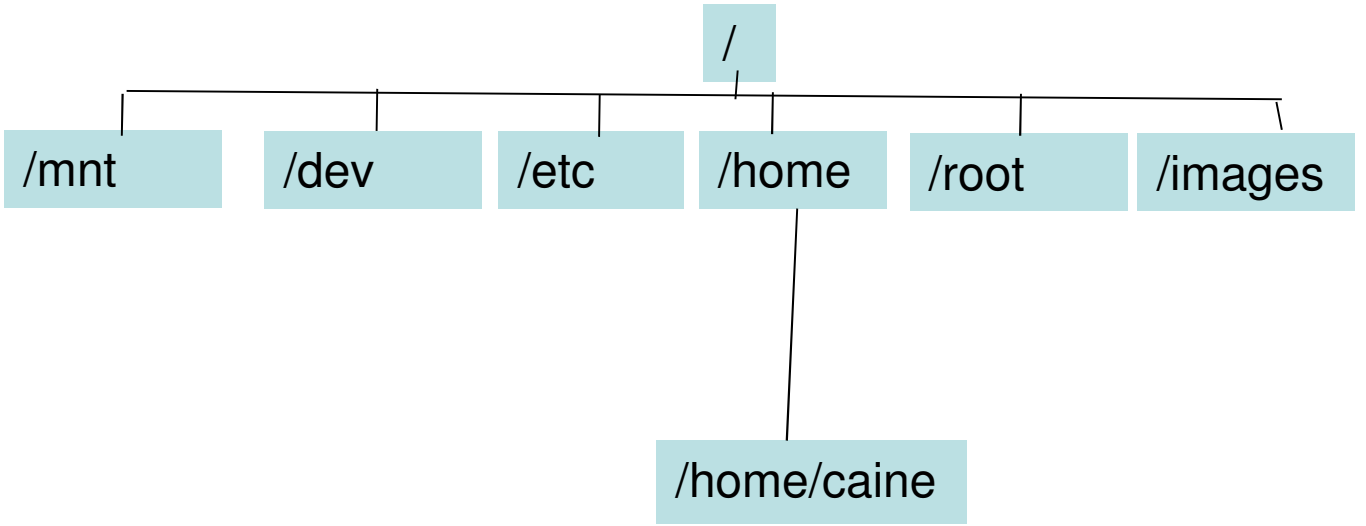

- This week we are accessing new directories in the directory tree
- /mnt – Empty directory useful for mounting images and disks
- /dev – Holds special files needed for mounting data
- /root – The administrator's HOME
- /images – A collection of disk images we are using for forensic tutorials.

Edinburgh Napier UNIVERSITY

sudo

- You run commands in Caine with the rights and privileges of the current user account. This is usually the user “caine”.
- Sometimes you need to run commands with more privileges..
- In Caine, this is done by prepending the command with the string “sudo”.

Edinburgh Napier UNIVERSITY

Example: sudo

$ ls -l /etc/shadow

```
-rw-r----- 1 root shadow 1209 2012-01-10 13:46 /etc/shadow
```

$ tail -1 /etc/shadow

```
tail: cannot open `/etc/shadow' for reading:
  Permission denied
```

$ sudo tail -1 /etc/shadow

```
telnetd:*:15349:0:99999:7:::
```

Reading binary files

- cat/more/less lets you look at text files.
- Sometimes you want to look at binary files, perhaps looking for normal english words which will help you understand the file.
- If you try this with cat you get rubbish:

$ cat /bin/cat

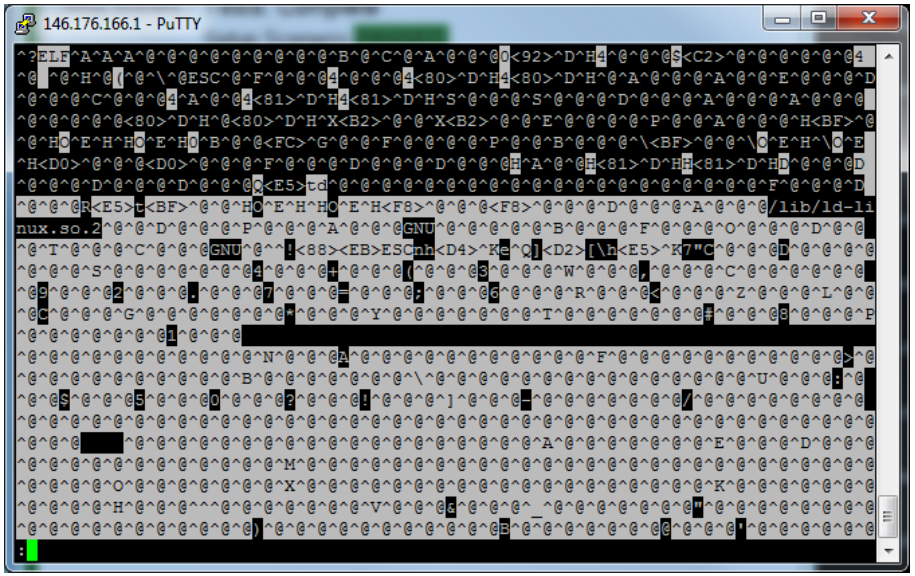

The strings command

- The “strings” command looks at binary files and tried to extract things which might be words.
- Really it just looks for normal characters which end in a NULL (\0) and which are over a certain length.
- Could be useful in understanding an unknown file:

$ strings /bin/cat

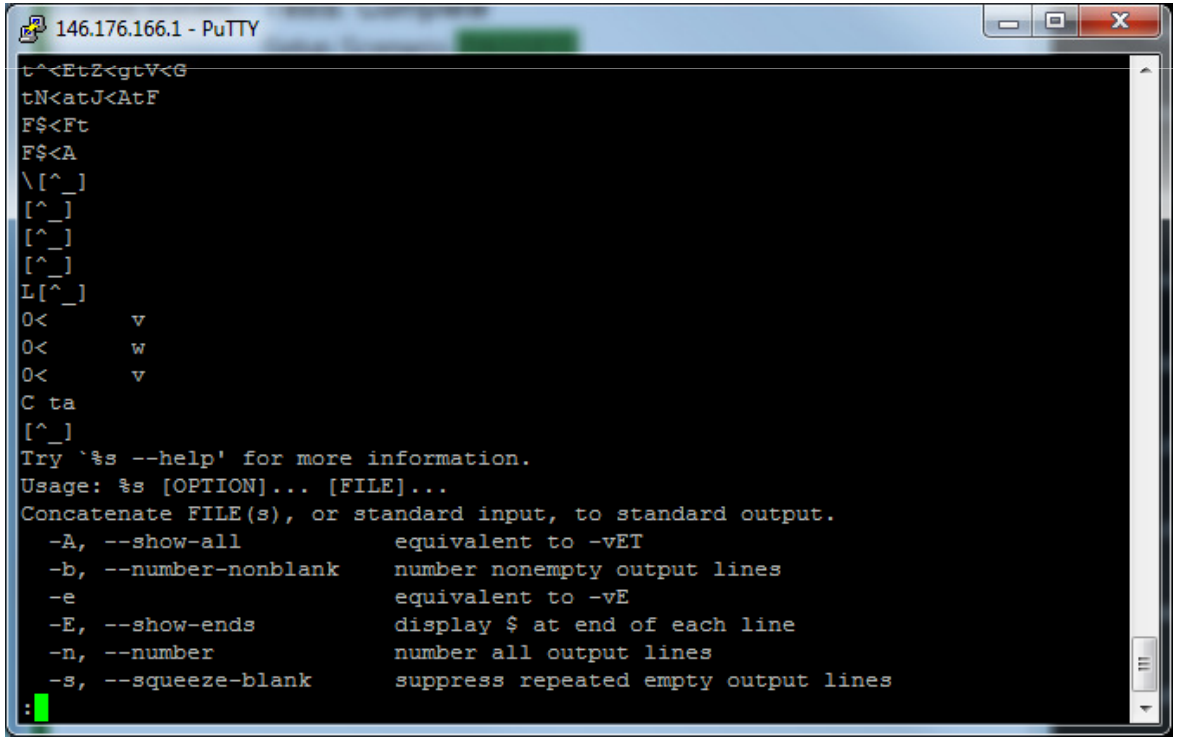

Edinburgh Napier UNIVERSITY

Disk Blocks

- Disk storage devices deal with data block transfers, not bytes.
- When a whole disk is copied to a file, it is called a disk image.
- When performing forensics, reading and writing blocks in disks or images may be necessary.
- In Linux, the system command to perform block reads and writes is called “dd”.

Edinburgh Napier UNIVERSITY

dd parameters

- dd takes many parameters. The ones we are interested in are:
  - if=filename – the filename or disk being read
  - of=filename – the filename or disk being written to
  - seek=blockno – Skip over blockno number of blocks from the beginning of the output file before starting to write
  - skip=blockno – Skip over blockno number of blocks from the beginning of the input file before starting to read
  - bs=512 – assume a block size of 512 bytes.
  - count=N – transfer N number of blocks
  - conv=notrunc – do not truncate the output file when finished.

dd example 1

$ dd if=file1.dd of=file2.dd bs=512

- The input file and the output file can be simple files, in which case the input file is copied to the output file. This is just like

  $ cp file1.dd file2.dd

- “if” or “of” could be actual disks. In Linux, disks can be accessed using special files in /dev. A hard drive could be “/dev/sdd” (which is sata disk 4).
- “if” or “of” can be other types of devices, some of which do special things.
- In Linux block size is usually 512 bytes. Always state “bs=512” unless you have reason not to.

Subset of blocks - skip

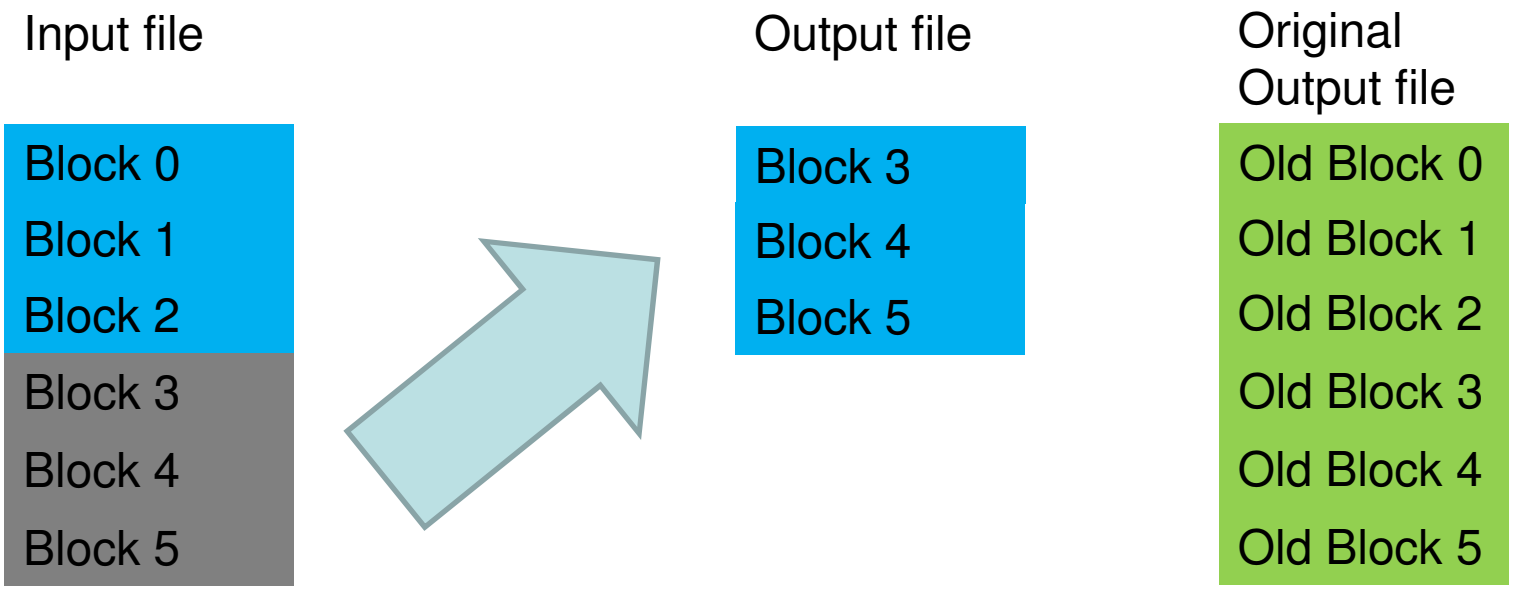

bs=512 count=3 **skip=3**

- Nothing remains of original destination.

Subset of blocks - seek

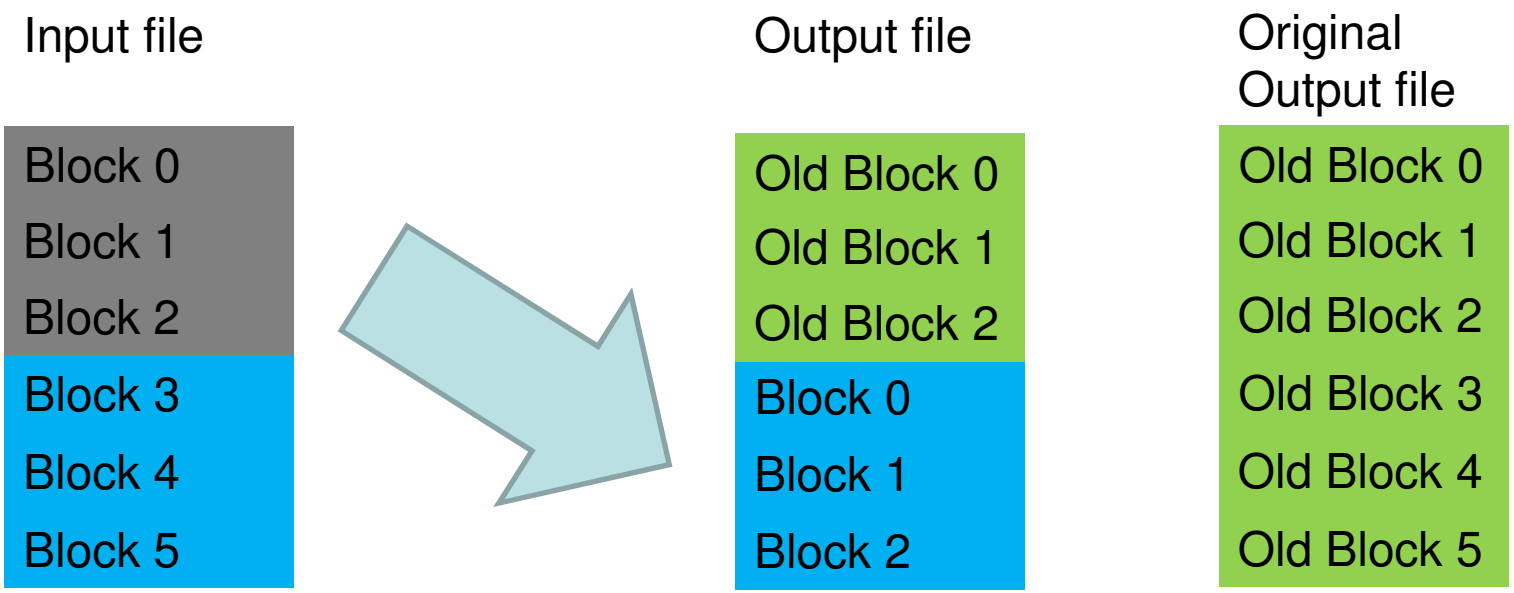

bs=512 count=3 **seek=3**

- Things in the destination BEFORE the seek point are kept

Seek and truncate

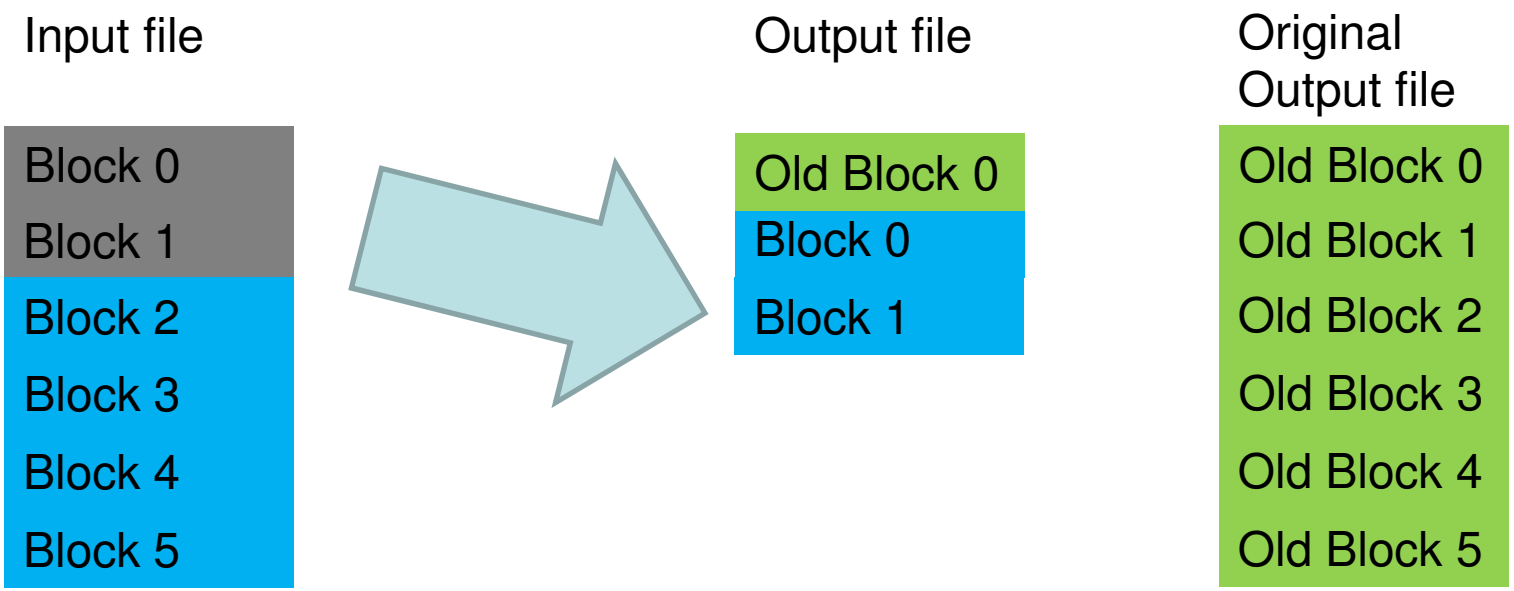

bs=512 count=2 seek=1

- Things in the destination AFTER the transfer are TRUNCATED

Seek and stop truncate

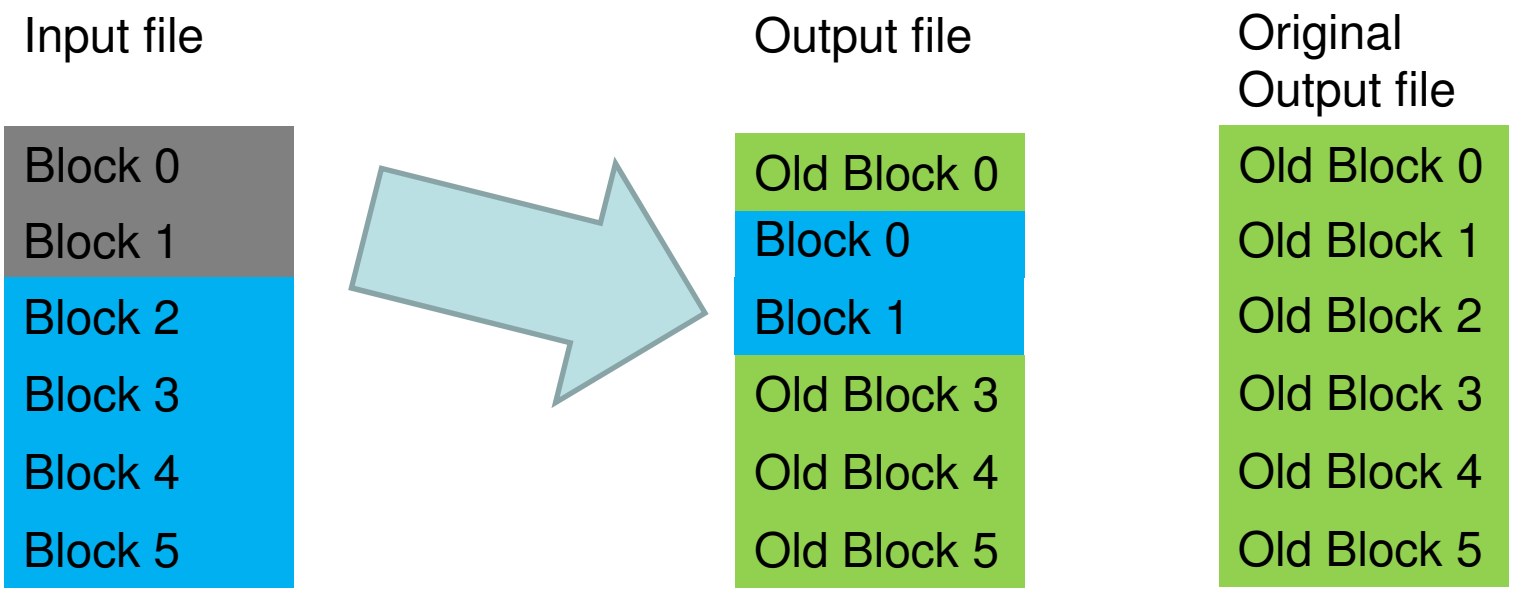

bs=512 count=2 seek=1 **conv=notrunc**

- Things in the destination not in the transfer are kept with notrunc.

Seek and Skip

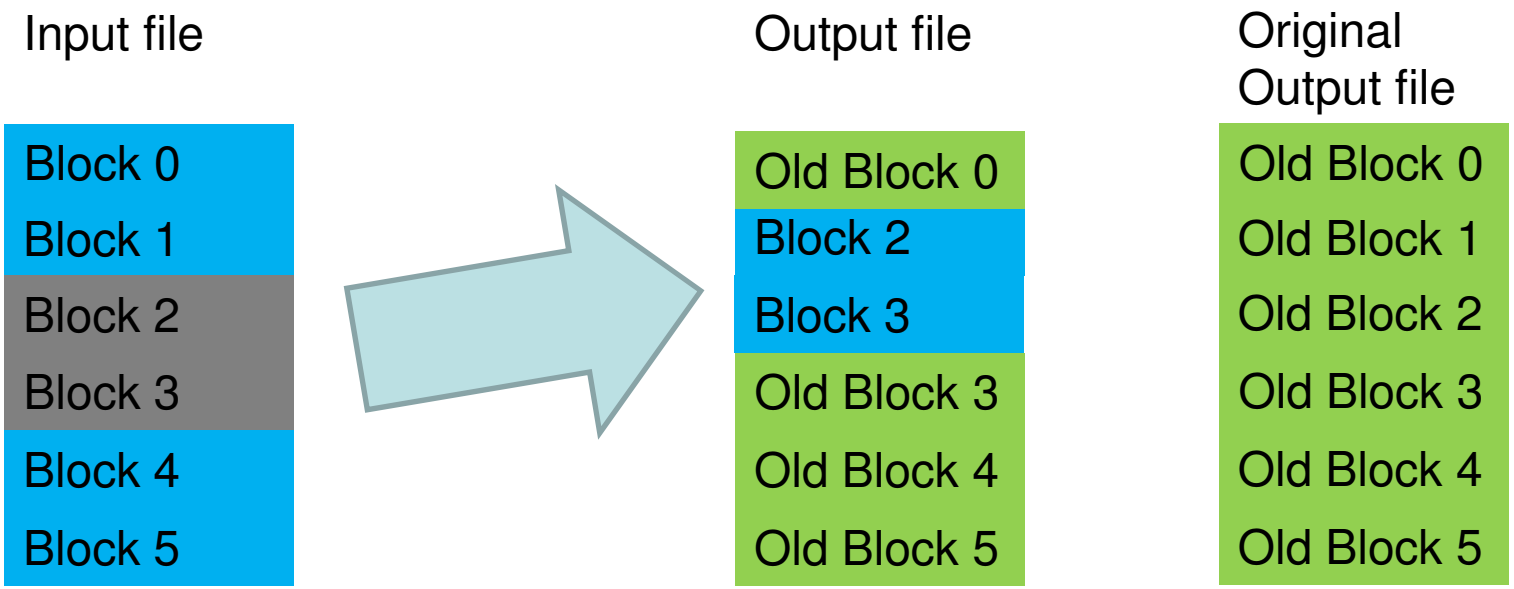

bs=512 count=2 **skip=2 seek=1** conv=notrunc

- Things in the destination not in the transfer are kept with notrunc.

Edinburgh Napier UNIVERSITY

Examples

- file1.dd contains 3 512 byte blocks. Write these to block 7 onwards of hard drive partition /dev/sda1. (Note real hard drives cannot be truncated)

$ dd if=file1.dd of=/dev/sda1 bs=512 count=3 seek=7 (conv=notrunc)

- Read block 9 from disk image file2.dd and write that block to file3.dd.

$ dd if=file2.dd of=file3.dd bs=512 count=1 skip=9

Edinburgh Napier UNIVERSITY

Questions:

- Copy block 11 and block 12 from file5.dd and save this as copy1.dd.

$ dd if=file5.dd of=copy1.dd bs=??? skip=??? count=???

- Copy block 55 from drive /dev/sdd and save this block as block 10 into a pre-existing disk image called image1.dd

$ dd if=/dev/sdd of=image1.dd bs=512 ????????????? count=1

Edinburgh Napier UNIVERSITY

Mount and umount

- If you have a hard drive on linux it has a device name.
- For instance in Caine the harddrive is called /dev/sda (SATA drive 1).
- Disks are partitioned up into disk partitions.
- In caine /dev/sda is partitioned into 1 partition. It is called /dev/sda1.
- If it had more partitions they would be /dev/sda1, /dev/sda3, etc.

- To access the contents of a partition the partition must be mounted at a mountpoint.
  - In Linux a mountpoint is simply an empty directory.
  - There is a commonly used empty directory called /mnt used for temporary mounts

```
$ mount /dev/sda1 /mnt
```

Understanding partition data

- To read a disk and find out its partition structure you might use any number of tools.
- An easy tool at the CLI is sfdisk. Remember to use sudo!

$ sfdisk –l –uS /dev/sda

```
Units = sectors of 512 bytes, counting from 0

   Device Boot    Start       End   #sectors  Id  System
/dev/sda1   *      2048   7167999    7165952  83  Linux
/dev/sda2             0         -          0   0  Empty
/dev/sda3             0         -          0   0  Empty
/dev/sda4             0         -          0   0  Empty
```

Edinburgh Napier UNIVERSITY

$ sfdisk –l –uS /dev/sda

```
Units = sectors of 512 bytes, counting from 0

   Device Boot    Start       End   #sectors  Id  System
/dev/sda1   *      2048   7167999    7165952  83  Linux
...
```

- “-l” shows you the partitions. This command can also set partitions. We will only be using “-l” for now.
- “-uS” says to report the sizes and other data in Sectors, which in this case is interchangable for 512 byte blocks.
- /dev/sda1 starts 2048*512 bytes into the disk, and is 7165952*512 bytes in size. It is a Linux partition (type 0x83).

Edinburgh Napier UNIVERSITY

Errors from sfdisk

- The Cylinders/Heads/Sectors of a disk (C/H/S) is given by the BIOS.
- This is needed to access the drive in C/H/S mode.
- C/H/S data may also be written to the hard drive, and may even be right...
- Some operating systems need this to be right.
- Linux does not care, and will access the disk in block mode.

- This happens a lot with disk images, as an image does not have a BIOS to tell you the geometry.
- Usually it is fine and you can ignore the issue...

Edinburgh Napier UNIVERSITY

```
Disk /dev/sda: 446 cylinders, 255 heads, 63 sectors/track
Warning: The partition table looks like it was made
  for C/H/S=*/48/49 (instead of 446/255/63).
For this listing I'll assume that geometry.
Units = sectors of 512 bytes, counting from 0

   Device Boot    Start       End   #sectors  Id  System
/dev/sda1   *      2048   7167999    7165952  83  Linux
               start: (c,h,s) expected (0,41,40) found (0,32,33)
               end: (c,h,s) expected (1023,47,49) found (446,47,49)
/dev/sda2             0         -          0   0  Empty
/dev/sda3             0         -          0   0  Empty
/dev/sda4             0         -          0   0  Empty
```

Edinburgh Napier UNIVERSITY

Mount parameters

- The first parameter of mount is the device name of the hard drive being mounted.
- The second parameter is the mountpoint.

$ mount /dev/sda1 /mnt

- Each partition must have been formatted already. In linux formats include ext3 and brfs. In Windows common format are fat/FAT16 and NTFS. This is usually detected automatically by Linux during the mount.

Edinburgh Napier UNIVERSITY

Removing a mount

- Once you are finished with the mount you can unmount it.

$ umount /dev/sda1

$ umount /mnt

- You can either unmount using the mountpoint name or the device name. Either one is fine.
- Note if you “cd /mnt” you cannot unmount /mnt. You will get an error that the drive is in use. Just “cd /” in that case before you umount.
- You can only mount devices, not files.

Mounting Files

- If you take a disk image (perhaps using dd) of a device you may want to mount that image.
- To do this you need to create a loop device. They are numbered /dev/loop0 to loop9.
  - This takes a disk image and gives it a device name
  - Remember to use sudo

```
$ losetup /dev/loop0 /images/imagename.dd
$ losetup -a
/dev/loop0: [0821]:13 (/images/imagename.dd)
$ mount /dev/loop0 /mnt
$ ls –l /mnt/
...
$ umount /dev/loop0
$ losetup –d /dev/loop0
```

Edinburgh Napier UNIVERSITY

Mounting Partitions in Files

- When you loop a disk image you have /dev access to the file as a disk, but no immediate access to the partitions individually.
- If /images/file1.dd is a whole disk, and you want to mount partition 1, you need to do some work!
- Firstly, you need to know the byte offset into the disk where the partition starts...
- So in the example of sfdisk above:

```
  Device     Boot     Start       End    #sectors   Id   System
/dev/sda1    *        2048    7167999     7165952   83   Linux
```

- As the start is 2048 blocks, the offset is 2048*512. Add this into the losetup using –a, e.g.

$ losetup /dev/loop3 /images/img1.dd –a 1048576

Edinburgh Napier UNIVERSITY

Snapshots

- Sometimes it would be nice to be able to take a disk image, and make changes to the image without the changes being permanent.
- For instance, if you have a 1TB disk image with a damaged partition table, you could copy the whole image and write to the copy.
  - But this would use up another 1TB of data.
- One alternative is to use snapshots, which allow you to change read-only data by pretending to write to the original data, but instead putting the writes into a separate file.
  - This new file holds all the changes, and so long as the number of changes are small this file should also be small.

Snapshot change

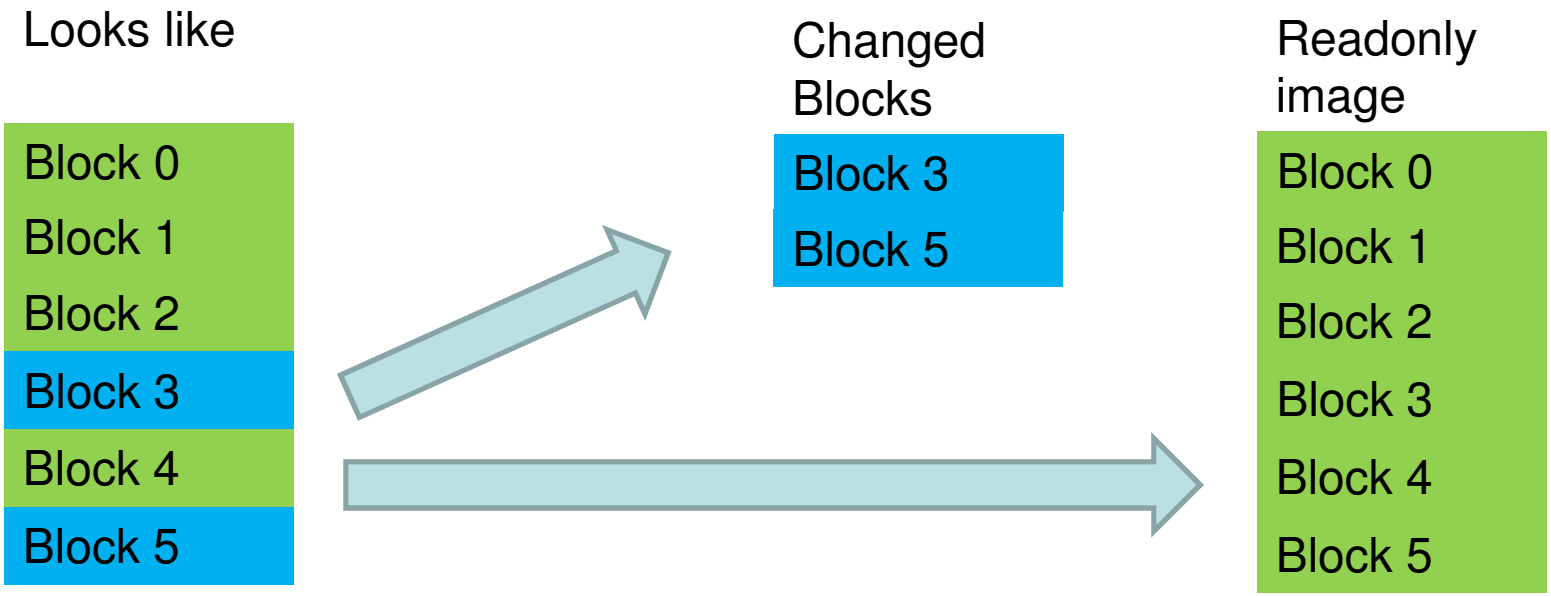

Build a snapshot device of /images/img1.dd

Edinburgh Napier UNIVERSITY

- File to hold the changes... Must be big enough!

$ dd if=/dev/zero of=/root/changes bs=512 seek=4095 count=1

- Build loop devices for the original and the changes

$ losetup /dev/loop5 /images/img1.dd (plus offset info for partitioning).

$ losetup /dev/loop6 /root/changes

- We need the blocksize of the original

$ blockdev –getsize /dev/loop5

- Create the snapshot

$ dmsetup create **sandbox** --table "0 THESIZE snapshot /dev/loop5 /dev/loop6 N 1“

- Mount the new device

$ mount /dev/mapper/**sandbox** /mnt

Edinburgh Napier UNIVERSITY

Remove shapshot

$ mount /mnt

$ dmsetup remove **sandbox**

$ losetup –d /dev/loop5

$ losetup –d /dev/loop6

Edinburgh Napier UNIVERSITY

Next Week

- My last lecture on Linux-focused commands
  - Although we will cover many more Caine commands for forensic-specific activities.
- Next week I will be looking at sorting, text-based diff, and understanding binary information in files.
- I will also be doing some preparation work for the first class test.

Edinburgh Napier UNIVERSITY

Assessment: Short-Answer Examples

- The short answer class test has no past papers yet (as this is a new module for this year).
- This section contains example questions which are of the same style as you might expect in the actual past paper.
- Obviously it is likely that the actual questions shown here are not the ACTUAL questions which will appear in the exam!
- Remember this short answer exam is CLOSED BOOK. You are not permitted to use the internet or access your notes during the exam.

Edinburgh Napier UNIVERSITY

Q1

- You have two disk images, “image1.dd” and “image2.dd”. Show the “dd” command to copy blocks 8,9, and 10 from image1.dd so that they appear as blocks 20,21,and 22 in image2.dd. All other blocks in image2.dd should be left unchanged.

Insert answer here:

Edinburgh Napier UNIVERSITY

Q2

- Consider the output of “sfdisk –l –uS /img/a.dd”.

```
 Device      Boot     Start       End    #sectors  Id  System
/img/a.dd1    *        4096    7167999     7165952  83  Linux
```

- Show the appropriate commands needed to mount the partition shown in /mnt.

Insert answer here:

Edinburgh Napier UNIVERSITY

Q3

- Briefly discuss two advantages in using snapshots during a forensic investigation.

Insert answer here: