Digital forensics

Andrej Brodnik

Operating system Unix

Chapter 18

- Development through history: *System V*, *HP-UX*, *BSD*, ...
- Open source versions that appeared later:
  - Linux: RedHat, SUSE, Ubuntu, ...
  - BSD: FreeBSD, OpenBSD, NetBSD

File system Hierarchy Standard

- *File system Hierarchy Standard* – FHS (http://www.pathname.com/fhs/pub/fhs-2.3.html)
- Linux Foundation took over the work (http://www.linuxfoundation.org/collaborate/workgroups/lsb/fhs)
- Mostly formalization of the BSD file system

Root directory

- */boot : Static files of the boot loader*
- */dev : Device files*
- */etc : Host-specific system configuration*
  - */etc/opt : Configuration files for /opt*
  - */etc/X11 : Configuration for the X Window System (optional)*
  - */etc/sgml : Configuration files for SGML (optional)*
  - */etc/xml : Configuration files for XML (optional)*
- /bin : Essential user command binaries (for use by all users)
- /sbin : System binaries
- /lib : Essential shared libraries and kernel modules
- /lib<qual> : Alternate format essential shared libraries (optional)

Root directory

- */home : User home directories (optional)*
- */root : Home directory for the root user (optional)*
- */media : Mount point for removable media*
- */mnt : Mount point for a temporarily mounted filesystem*
- */opt : Add-on application software packages*
- */srv : Data for services provided by this system*
- */tmp : Temporary files*
- */usr, /var : Separate hierarchies*

/usr directory

- Contains read-only files
- Used simultaneously by different systems
- Doesn’t contain files that are specific to a particular system
- Exception: /usr/local, which is the local directory of a particular system

/var directory

- Contains files that change over time
  - Postal and print queues
  - Logging
  - Data (databases etc)
  - Temporary files

System files

- Operating system is designed so that system files are user-friendly ➔ regular text files
  - Configuration files: hosts, syslog.conf, …
    - Usually in the directory etc (/etc, /usr/local/etc, /opt/etc, …)
  - Logging : mail, cups, …
    - Usually in the directory log (/var/log, /usr/local/var/log, /opt/var/log)

Configuration files

```
# $FreeBSD: release/9.0.0/etc/snmpd.config 216595 2010-12-20 17:28:15Z syrinx $
#
# Example configuration file for bsnmpd(1).
#

#
# Set some common variables
#
location := "Room 200"
contact := "sysmeister@example.com"
system := 1     # FreeBSD
traphost := localhost
trapport := 162

#
# Set the SNMP engine ID.
#
# The snmpEngineID object required from the SNMPv3 Framework. If not explicitly set via
# this configuration file, an ID is assigned based on the value of the
# kern.hostid variable
# engine := 0x80:0x10:0x08:0x10:0x80:0x25
# snmpEngineID = $(engine)
```

Logging

Mar 8 00:00:00 svarun newsyslog[85254]: logfile turned over
Mar 8 00:00:12 svarun postfix/smtpd[85247]: connect from S0106c0c1c0ddffcf.vf.shawcable.net[70.69.32.154]
Mar 8 00:00:12 svarun postfix/smtpd[85247]: NOQUEUE: reject: RCPT from S0106c0c1c0ddffcf.vf.shawcable.net[70.69.32.154]: 554 5.7.1 Service unavailable; Client host [70.69.32.154] blocked using bl.spamcop.net; Blocked - see http://www.spamcop.net/bl.shtml?70.69.32.154; from=<unscrupulousnessiw2@deltamar.net> to=<xxxx@brodnik.org> proto=ESMTP helo=<deltamar.net>
Mar 8 00:00:12 svarun postfix/smtpd[85247]: lost connection after DATA from S0106c0c1c0ddffcf.vf.shawcable.net[70.69.32.154]

Data storage and hiding

- Simplified organization of the disk with FAT file system

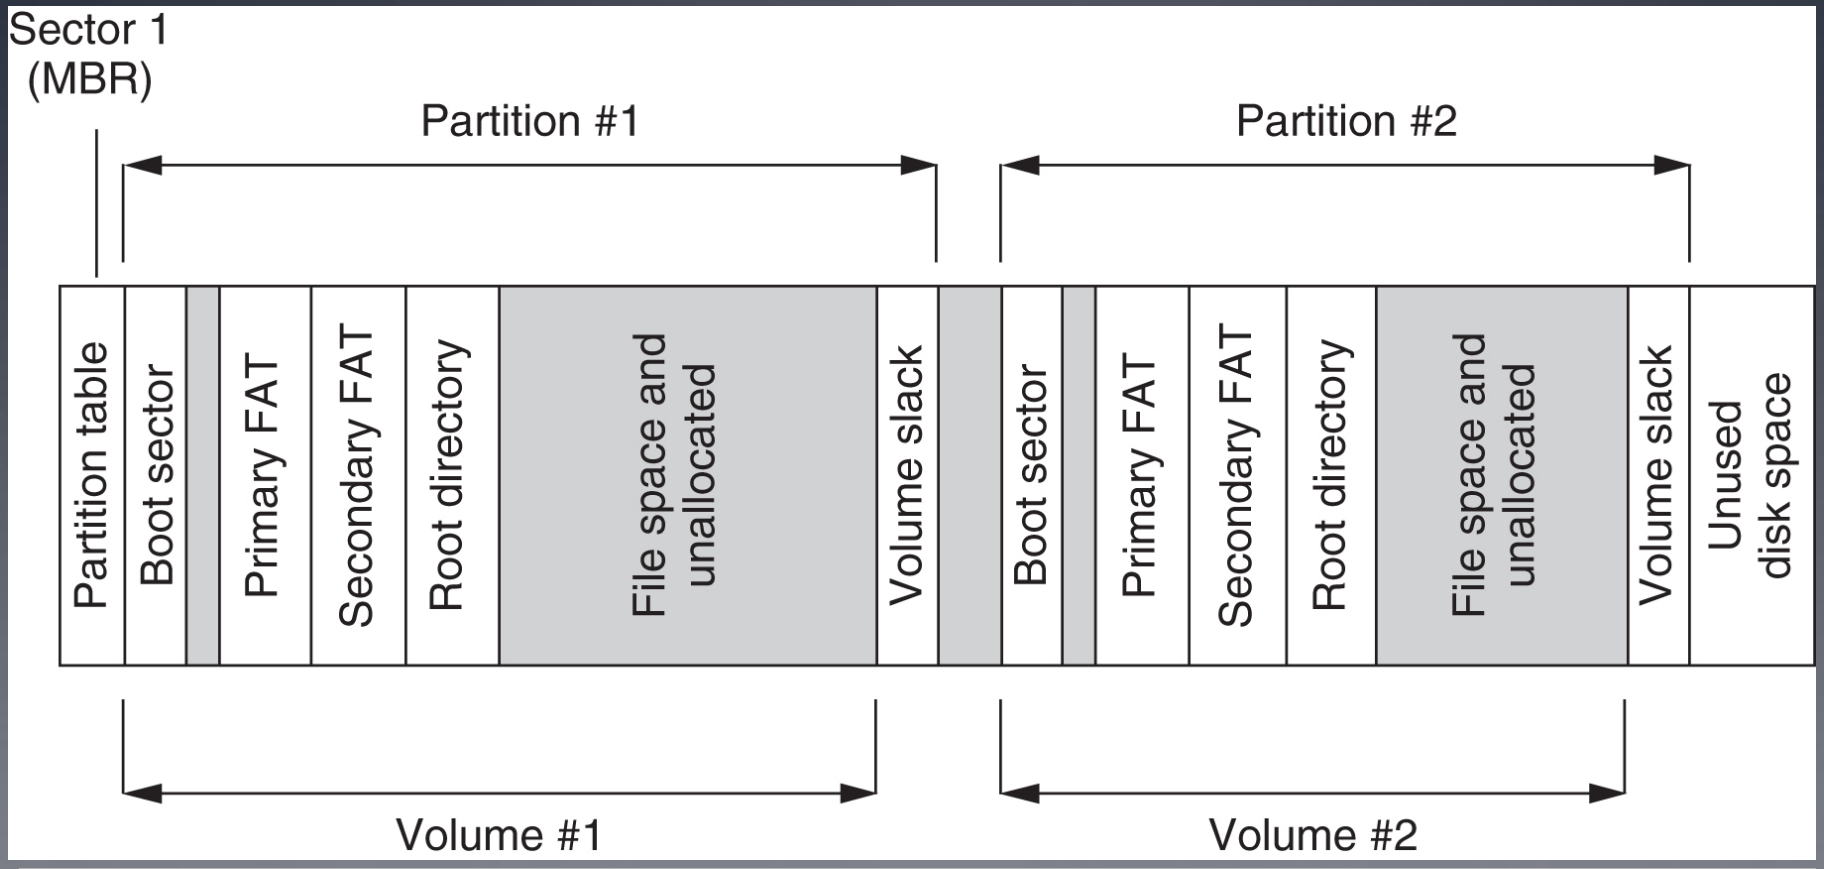

File systems

- We have directories and index nodes(*inode*)
- Inode has a similar function to FAT and MFT at the same time
- Directory is just a special file type
  - We have more special files: links, pipe, socket …

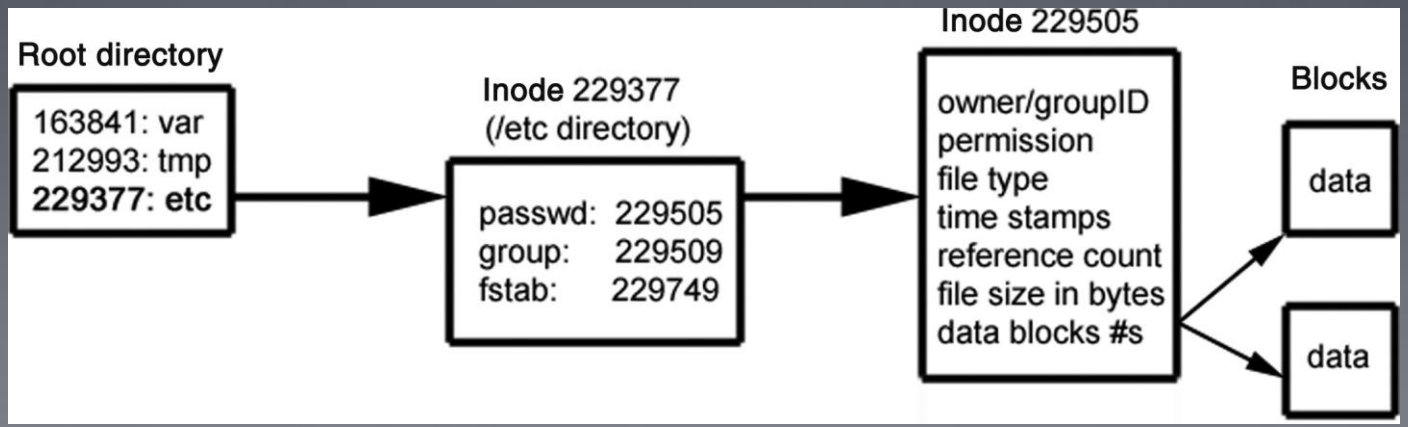

File systems

- The oldest: Unix File System – UFS
- More recent and used in Linux: ext2 in ext3
  - There are also ext and ext4
- There are a number of other file systems

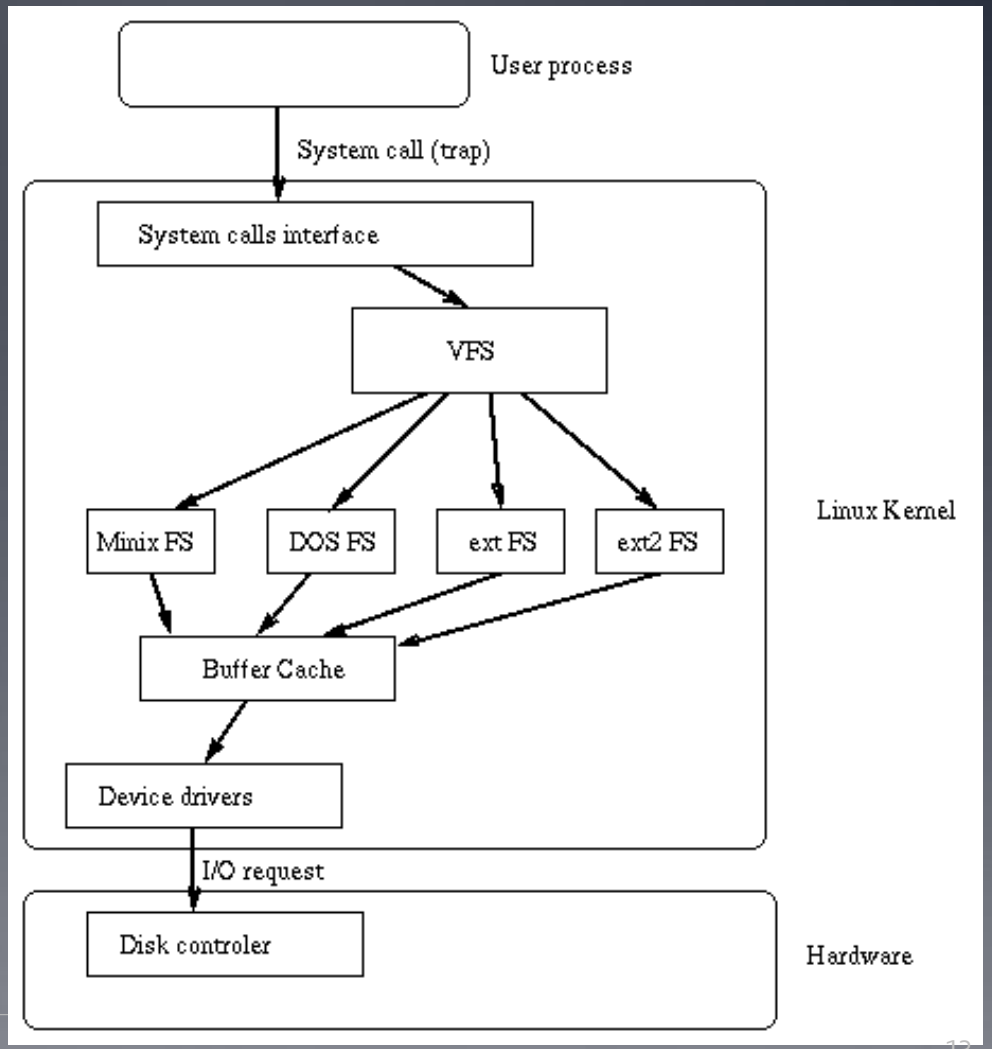

Time in the Unix operating system

- Time is measured in seconds
- Stored as a number, which begins on 1st of December 1970
  - If time is stored as a 32-bit number, there will be an overflow on Tuesday, December 19th 2038 at 03:14:07 UTC – Y2K38 problem
- UTC – *Coordinate Universal Time*: a harmonized definition of time that takes into account leap years, leap seconds, …
  - The last leap second occurred on 31st January 2016
  - harmonized time between several atomic hours
  - one of the successors of GMT

UFS file systems

- Defined when VFS was introduced in BSD4.2
- Used in *BSD systems
- Later used in Solaris OS

vir: *Solaris Internals*, *The UFS File System*, Updated by Frank Batschulat, Shawn Debnath, Sarah Jelinek, Dworkin Muller, and Karen Rochford

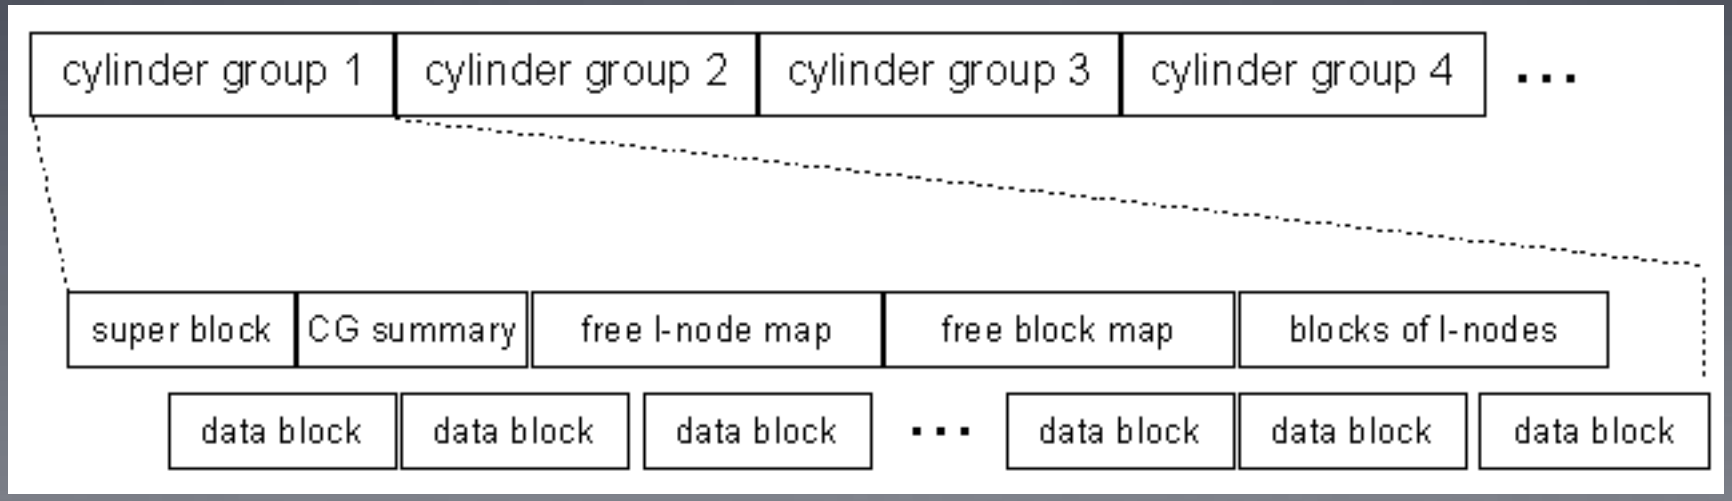

UFS – index node

```
struct dinode {
  u_int16_t       di_mode;        /*   0: IFMT, permissions; see below. */
  int16_t         di_nlink;       /*   2: File link count. */
  union {
    u_int16_t oldids[2];          /*   4: Ffs: old user and group ids. */
    int32_t   inumber;            /*   4: Lfs: inode number. */
  } di_u;
  u_int64_t       di_size;        /*   8: File byte count. */
  int32_t         di_atime;       /*  16: Last access time. */
  int32_t         di_atimensec;   /*  20: Last access time. */
  int32_t         di_mtime;       /*  24: Last modified time. */
  int32_t         di_mtimensec;   /*  28: Last modified time. */
  int32_t         di_ctime;       /*  32: Last inode change time. */
  int32_t         di_ctimensec;   /*  36: Last inode change time. */
  ufs_daddr_t     di_db[NDADDR];  /*  40: Direct disk blocks. */
  ufs_daddr_t     di_ib[NIADDR];  /*  88: Indirect disk blocks. */
  u_int32_t       di_flags;       /* 100: Status flags (chflags). */
  int32_t         di_blocks;      /* 104: Blocks actually held. */
  int32_t         di_gen;         /* 108: Generation number. */
  u_int32_t       di_uid;         /* 112: File owner. */
  u_int32_t       di_gid;         /* 116: File group. */
  int32_t         di_spare[2];    /* 120: Reserved; currently unused */
}
```

ufs/dinode.h

UFS – file systems

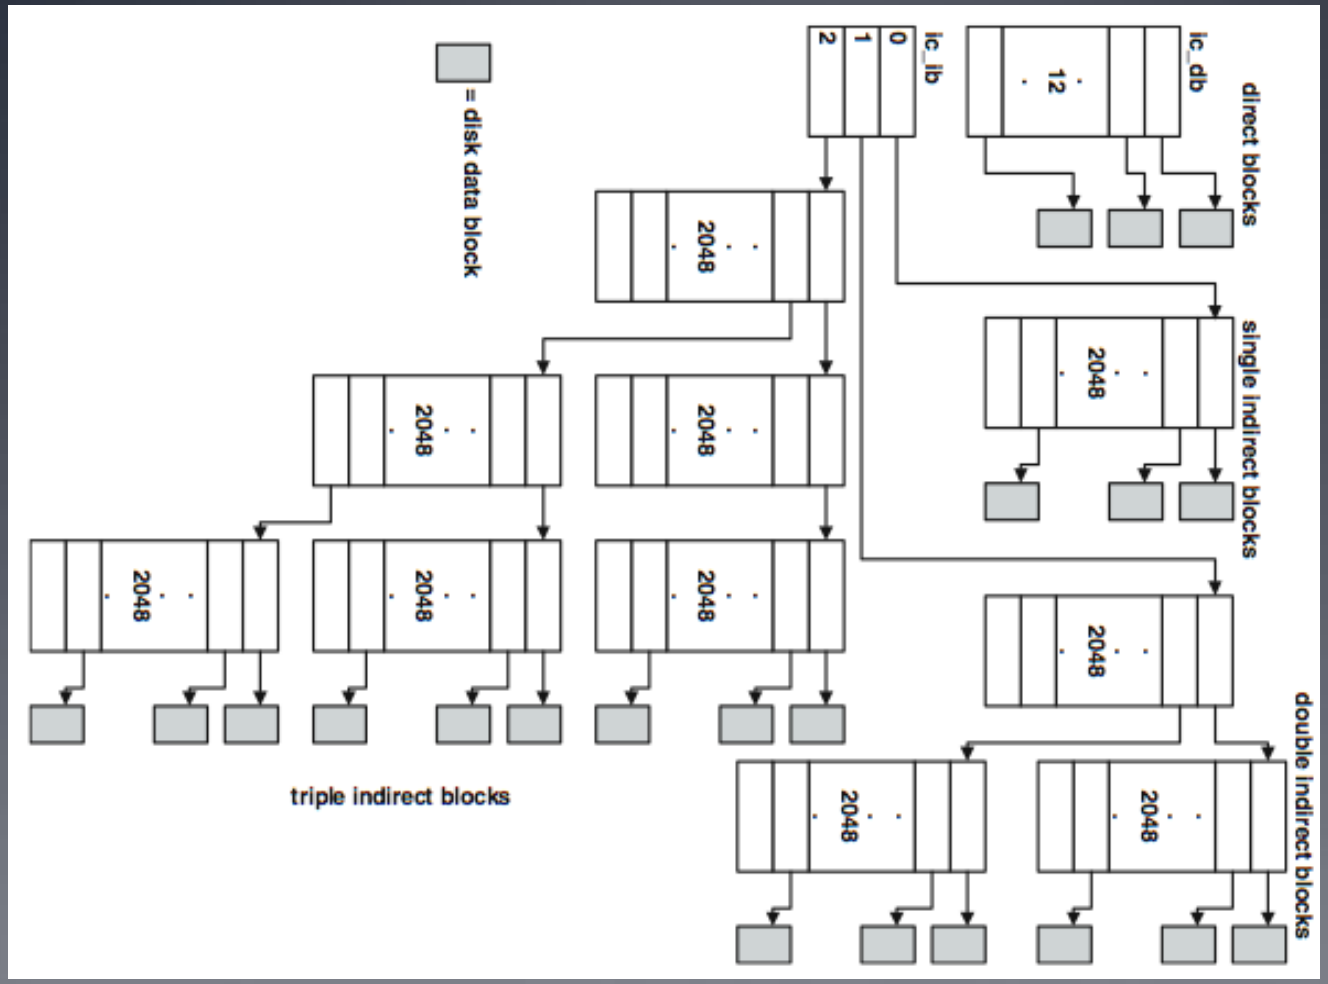

UFS – directory file

- a special file that consists of parts of the directory
- System V had a predefined file size
- The root directory is described in inode 2
- Each directory has a special entry that notes where its parent is

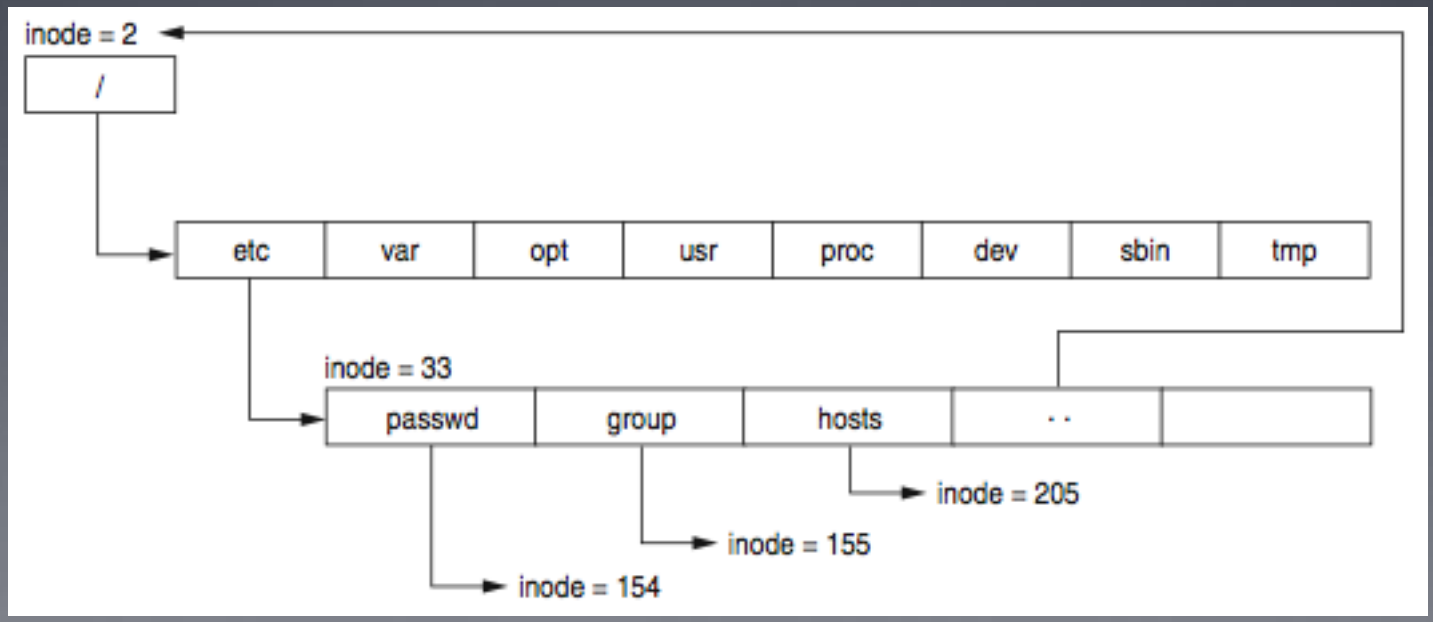

UFS – directory entry

```
#define MAXNAMLEN 255
struct direct {
  u_int32_t d_ino;       /* inode number of entry */
  u_int16_t d_reclen;    /* length of this record */
  u_int8_t  d_type;      /* file type, see below */
  u_int8_t  d_namlen;    /* length of string in d_name */
  char      d_name[MAXNAMLEN + 1];
                         /* name with length <= MAXNAMLEN */
};
```

ufs/dir.h

- *Challenge:* what is the record reclen intended for? Can this be used to hide data?
- *Challenge:* what is ACL? How is it implemented in UFS?

UFS – superblock

- Superblock stores the description of the cylinder group's configuration
- Scattered around the disc - at the beginning of each group of cylinders
- To save the configuration – if one record is lost
- **`dumpfs`** tool
- *Challenge:* find the structure of superblock. How do we know that we are dealing with the UFS file system? Where is it written? Read the superblock from your unix file system and find out for which file system it is.

File system ext2

- The basic structure is similar to that of UFS
- Instead of groups of cylinders, we are talking about block groups
- Directories and index nodes - like in UFS

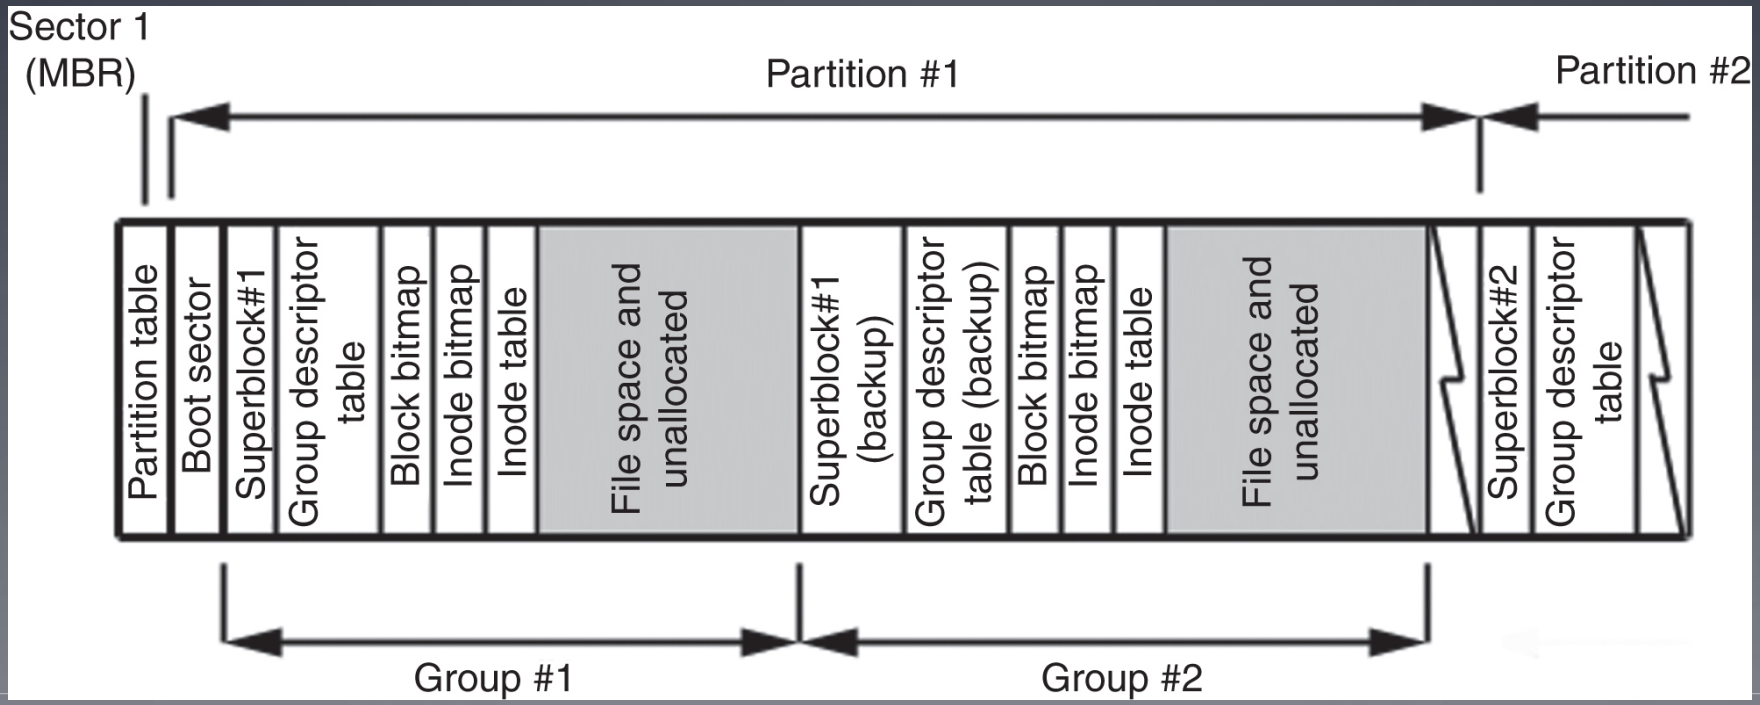

File system ext2

- Tool for viewing disks: Linux Disk Editor (LDE) (http://lde.sourceforge.net/)

```
                        lde v2.6.0 : ext2 : /dev/hdd2
Inode:          2 (0x00000002)  Block:           0 (0x00000000)

0x00000002: drwxr-xr-x  21      4096 .
0x00000002: drwxr-xr-x  21      4096 ..
0x0000000B: drwxr-xr-x   2     16384 lost+found
0x00008001: drwxr-xr-x   2      4096 boot
0x00010001: drwxr-xr-x  17     77824 dev
0x00020001: drwxr-xr-x   2      4096 proc
0x0000000C: -rw-r--r--   1         0 .autofsck
0x00028001: drwxr-xr-x  17      4096 var
0x00034001: drwxrwxrwt   8      4096 tmp
0x00038001: drwxr-xr-x  49      4096 etc
0x00048001: drwxr-xr-x  15      4096 usr
0x00598003: drwxr-xr-x   2      4096 bin
0x00640003: drwxr-xr-x   3      4096 home
0x0064C003: drwxr-xr-x   2      4096 initrd
0x00650003: drwxr-xr-x   7      4096 lib
0x00660003: drwxr-xr-x   4      4096 mnt
0x0066C003: drwxr-xr-x   2      4096 opt
0x00670003: drwxr-x---   7      4096 root
0x0067C003: drwxr-xr-x   2      4096 sbin
0x0044C04C: drwxr-xr-x   2      4096 misc
0x000E0021: drwxr-xr-x   4      4096 e1
```

File system ext2

```
                        lde v2.6.0 : ext2 : /dev/hdd2
Inode:     229505 (0x00038081)  Block:          0 (0x00000000)

-rw-r--r--   1 root     root         1186  Tue Sep 24 08:57:40 2002

TYPE: regular file  LINKS:   1              DIRECT BLOCKS=0x000703F9
MODE: \0644         FLAGS: \10
UID: 00000(root)    GID: 00000(root)
SIZE: 1186          SIZE(BLKS): 8

ACCESS TIME:        Tue Nov 26 11:10:18 2002
CREATION TIME:      Tue Sep 24 08:57:40 2002
MODIFICATION TIME:  Tue Sep 24 08:57:40 2002
DELETION TIME:      Wed Dec 31 19:00:00 1969




                                                  INDIRECT BLOCK=
                                                  2x INDIRECT BLOCK=
                                                  3x INDIRECT BLOCK=
```

ext2 – index nodes

```
struct ext2_inode {
    __u16   i_mode;         /*   0: File mode */
    __u16   i_uid;          /*   2: Owner Uid */
    __u32   i_size;         /*   4: Size in bytes */
    __u32   i_atime;        /*   8: Access time */
    __u32   i_ctime;        /*  12: Creation time */
    __u32   i_mtime;        /*  16: Modification time */
    __u32   i_dtime;        /*  20: Deletion Time */
    __u16   i_gid;          /*  24: Group Id */
    __u16   i_links_count;  /*  26: Links count */
    __u32   i_blocks;       /*  28: Blocks count */
    __u32   i_flags;        /*  32: File flags */
    __u32   l_i_reserved1;  /*  36: OS dependent 1 */
    __u32   i_block[EXT2_N_BLOCKS];/* 40: Pointers to blocks */
    __u32   i_generation;   /* 100: File version (for NFS) */
    __u32   i_file_acl;     /* 104: File ACL */
    __u32   i_dir_acl;      /* 108: Directory ACL */
    __u32   i_faddr;        /* 112:  Fragment address */
    __u8    l_i_frag;       /* 116:  Fragment number */
    __u8    l_i_fsize;      /* 117:   Fragment size */
    __u16   i_pad1;         /* 118:   */
    __u32   l_i_reserved2[2]; /* 120: OS dependent 2 */
};
```

ext2fs/ext2_fs.h

ext2 – index nodes

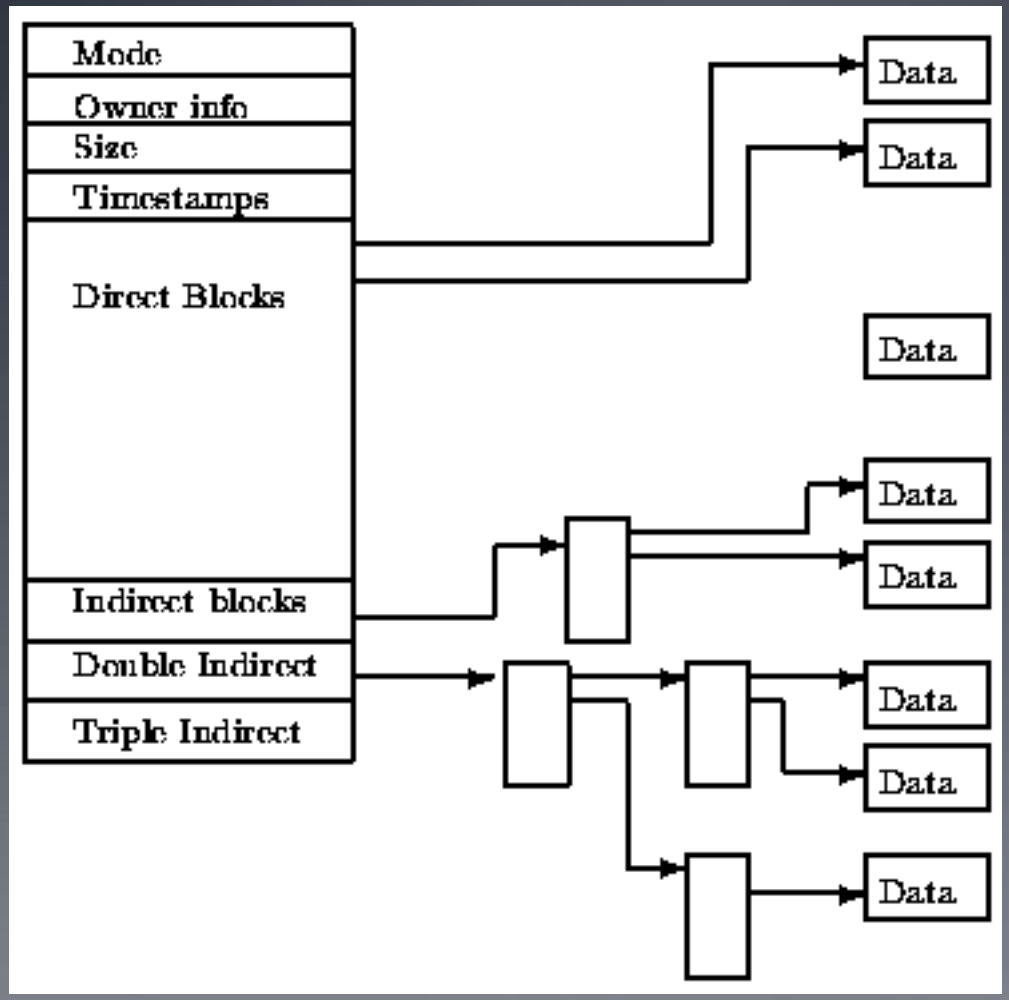

Directory file

- A special file that consists of parts of the directory
- System V had a predefined file size
- The root directory is described in inode 2
- Each directory has a special entry .. that tells where the parent is

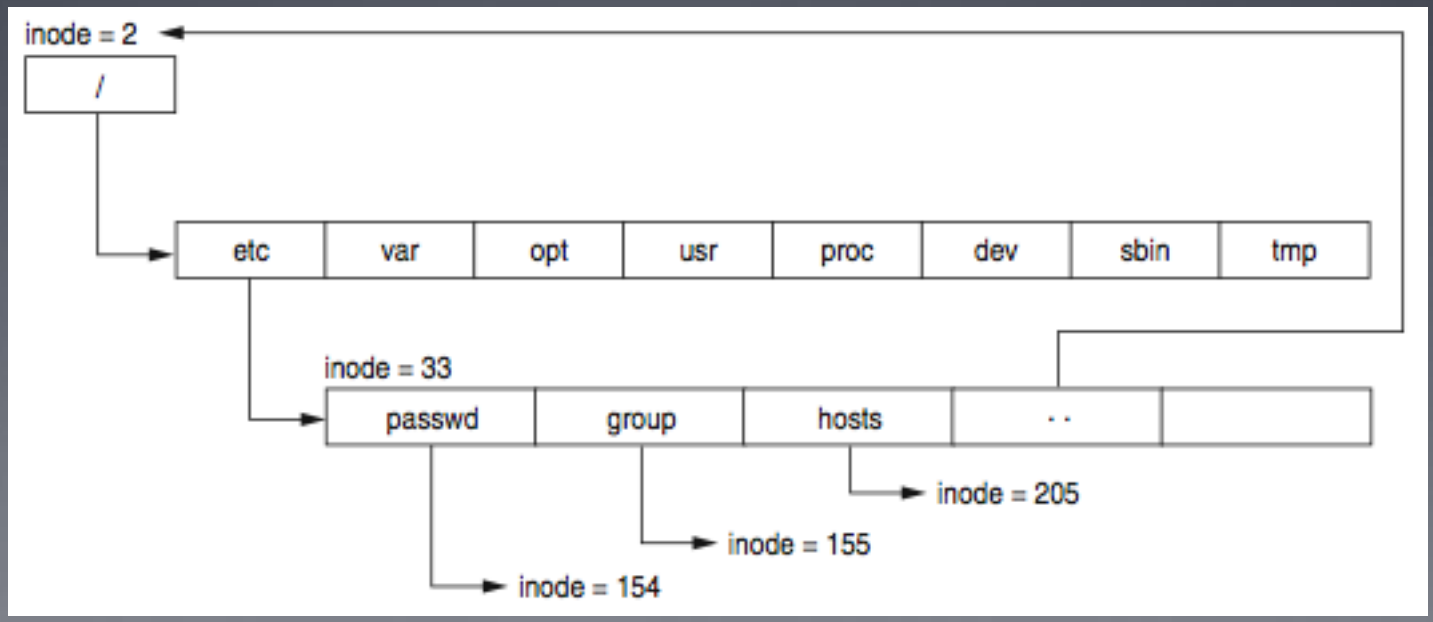

ext2 – Directory entry

```
#define EXT2FS_MAXNAMLEN      255
struct  ext2fs_direct {
  u_int32_t e2d_ino;       /* inode number of entry */
  u_int16_t e2d_reclen;    /* length of this record */
  u_int8_t e2d_namlen;     /* length of string in d_name */
  u_int8_t e2d_type;       /* file type */
  char     e2d_name[EXT2FS_MAXNAMLEN];
                           /* name with length <=
   EXT2FS_MAXNAMLEN */
};
```

ext2fs/ext2fs_dir.h

ext2 – superblock

- The superblock stores the description of the block group configuration
- Scattered across the disk - at the beginning of each block of blocks
  - to save the configuration if one record is lost
- Tool **`dumpfs`**

- *Izziv:* poiščite strukturo nadbloka ext2. Primerjajte jo s strukturo UFS superbloka.

File system ext3

- Author Stephen Tweedie 1999 / 2000 / 2001
- The basic structure is the same as for the ext2 file system
  - Split into blocks of blocks including a superblock
  - Directories and index nodes
  - Keeping track of the disk
- The option of saving the log structure is added
- The basic OS Linux file system

Journals ext3

- Journals contain records of all changes to the file system
- Journal's structure allows for three types of journals:
  - comprehensive journal: saves everything; both metadata and content - the most secure
  - ordered: only metadata is stored but only after a successful operation - medium safe
  - writeback: similar to sequential, saving log records at the same time as actual records - least secure

Journals ext3

- Journal is a sequential file
- Records are stored in front of the first group of blocks
- The log group is similar to the block group:
  - Journal superblock
  - Transaction descriptions

Journal ext3

- The transaction description contains three types of blocks:
  - Descriptor block: start of a transaction
  - Metadata blocks: transaction descriptions
  - Commit block: completion of the transaction
  - Revoke block: if an error occurs and contains a list of blocks in the file system that need to be reinstalled (restored)
- All (including superblock) start with a magical number:

  JFS_DESCRIPTOR_BLOCK 1

  JFS_COMMIT_BLOCK 2

  JFS_SUPERBLOCK_V1 3

  JFS_SUPERBLOCK_V2 4

  JFS_REVOKE_BLOCK 5

Journal ext3

- *Challenge:* Consider the structure of a superblock (e.g. http://linuxsoftware.co.nz/wiki/ext3) . Get a block from your file system and comment on its contents.
- *Challenge:* How do we restored a deleted file in ext2 and in ext3? What about in UFS?

File systems

- There are other file systems:
  - reiserFS, XFS, gfs, afs, ext4, HSM, …

- *Challenge:* Make a similar analysis for the mentioned systems as we did for UFS and ext.
- *Challenge:* Compare the described file systems – in which can we hide data?
- *Challenge:* Prepare a file system for your colleague and he must figure out which one it is.

Forensic sources

- We use stand-alone operating systems to analyse the disc image
- Example: Helix (Ubuntu)

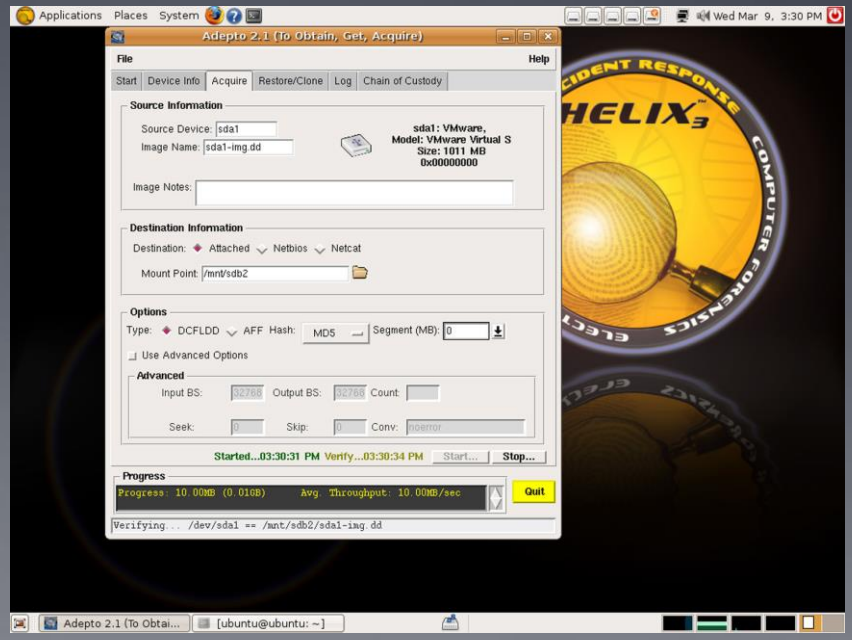

- *Challenge:* Prepare the Helix CD and check what tools are already on it.
- *Challenge:* Find some other similar systems.

Forensic sources

- Tool *SleuthKit* with *Autopsy Forensic Browser*

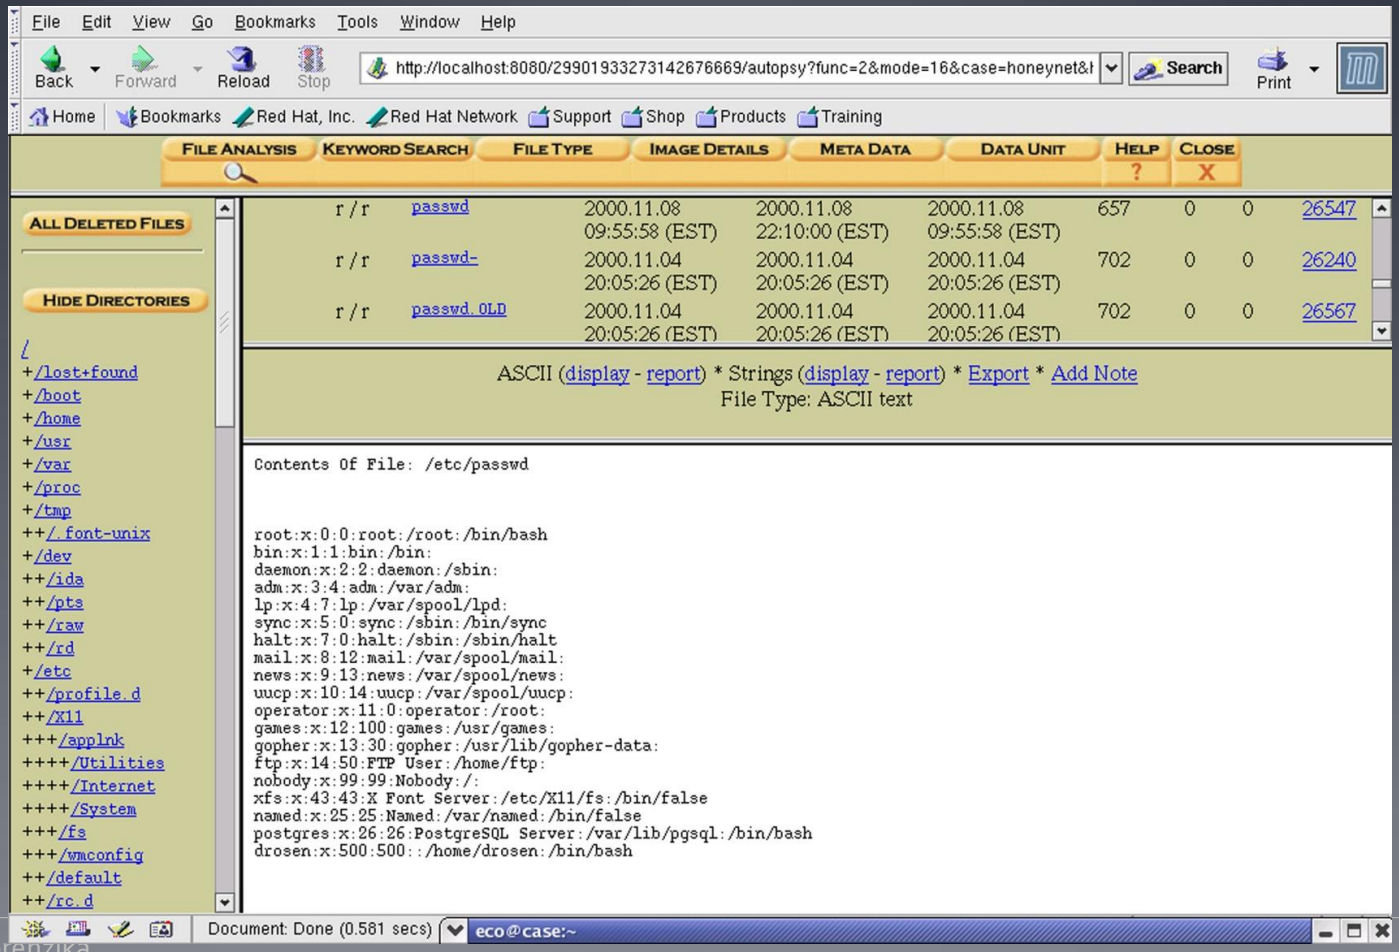

Forensic sources – research with *SleuthKit*

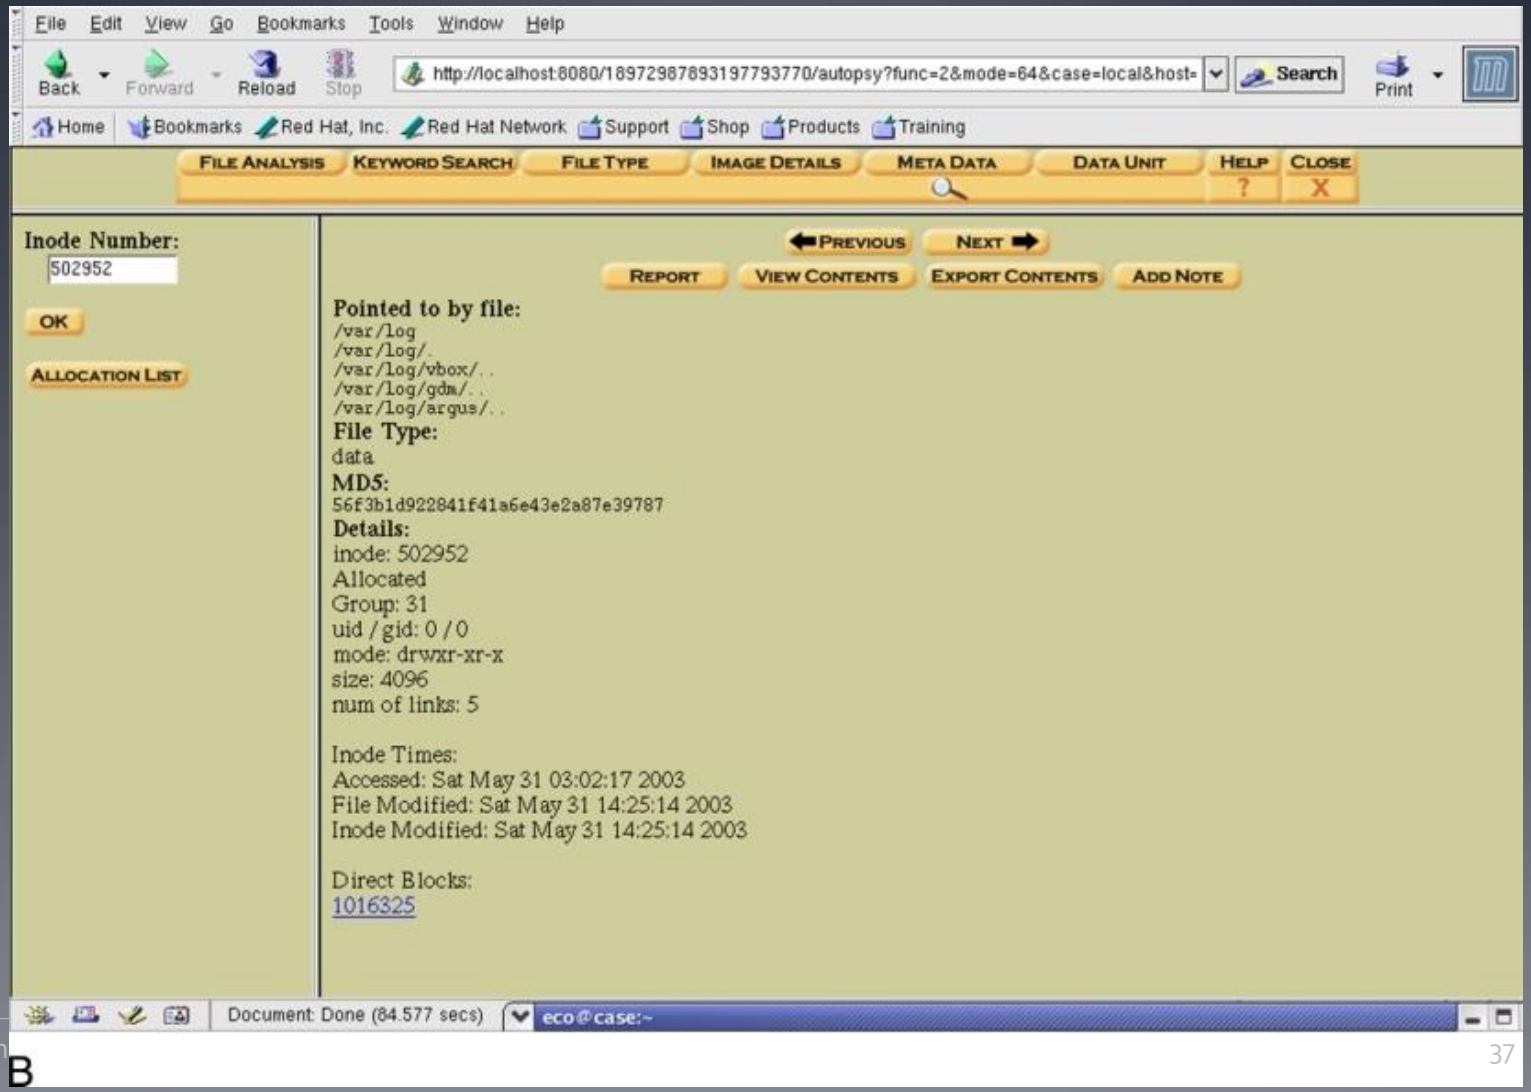

Forensic sources – research with *SleuthKit*

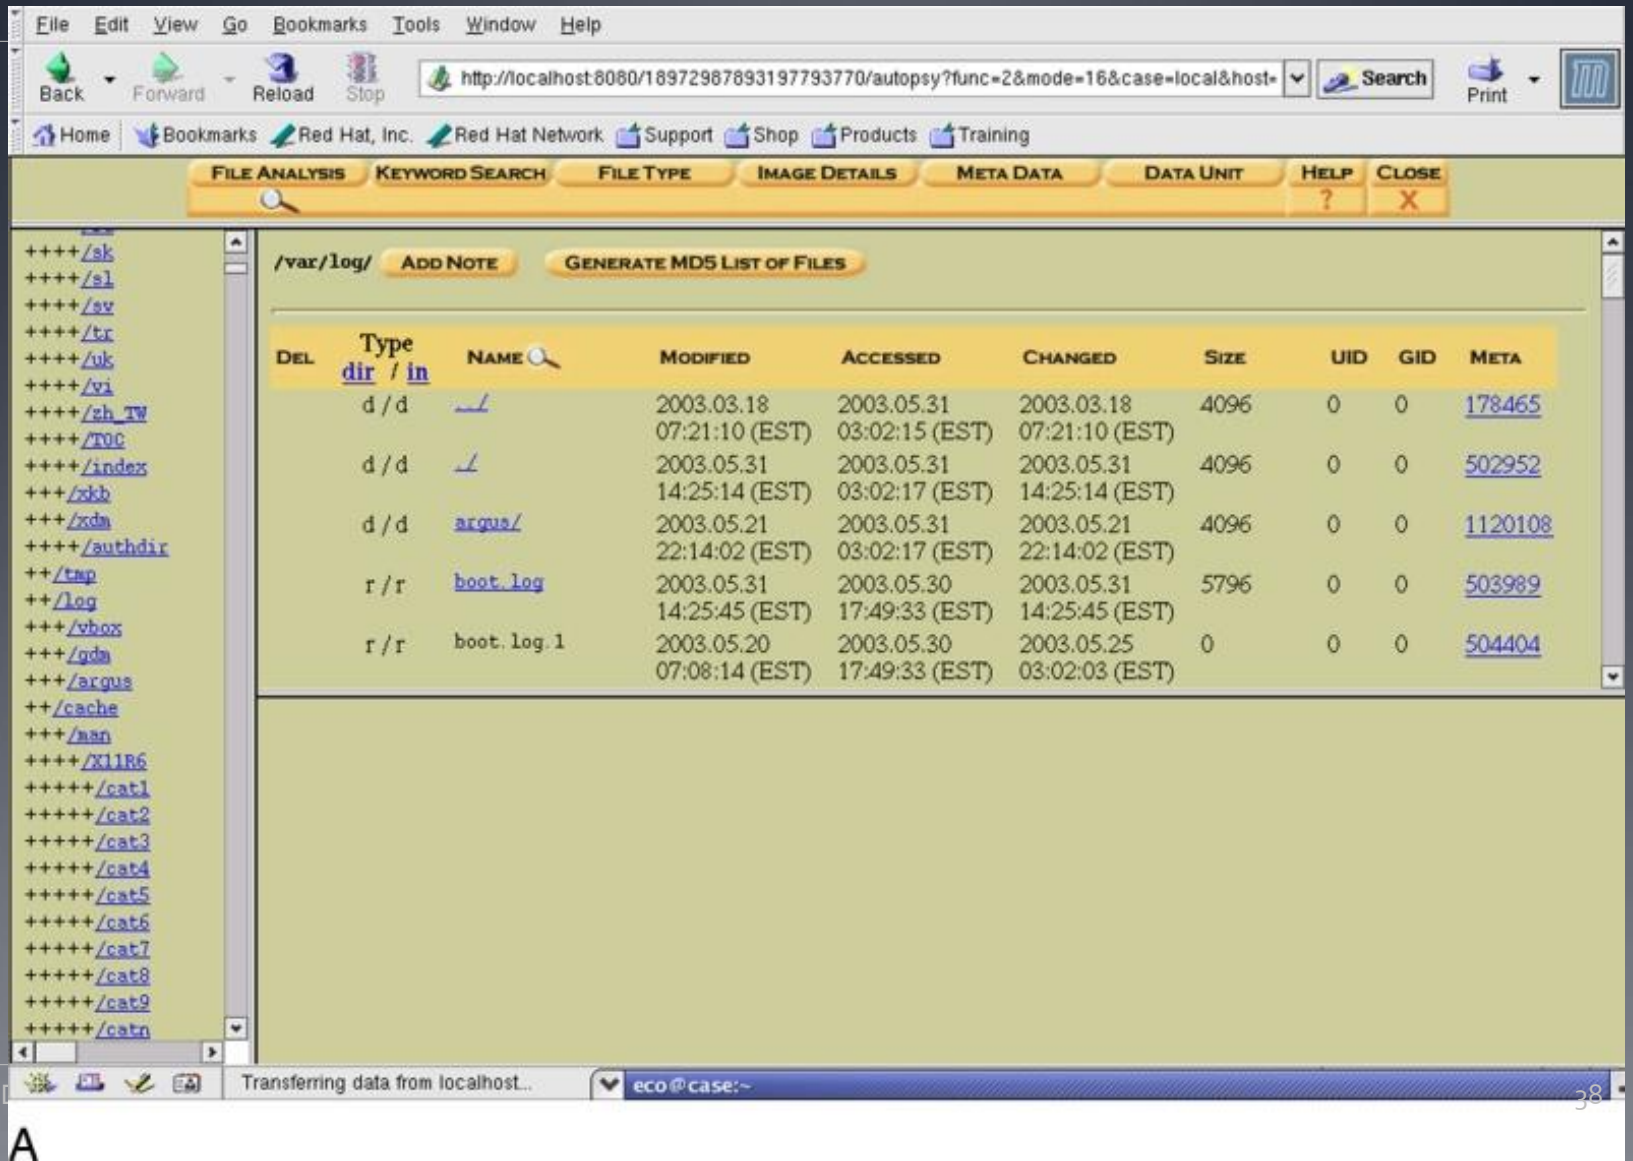

Forensic sources

- Video *File System Forensic Analysis* (*www.youtube.com/watch?v=rmG8yt1WpuA*)
- Different organizations
  - SANS Institute (*Sysadmin, Audit, Networking, and Security*): courses, literature…
  - The Honeynet Project (http://www.honeynet.org/)

- *Challenge:* Check out challenges on http://www.honeynet.org/challenges and try them.

Forensic sources

- Some interesting and rich references:
  - B. Carter, *File system forensic analysis*. Addison-Wesley, 2005.
  - Gregorio Narváez, *Taking advantage of Ext3 journaling file system in a forensic investigation*. SANS Institute, 2007.