

PV204: Disk encryption

Laboratory of Security and Applied Cryptography (LaBAK)

May 12, 2014, xbroz@fi.muni.cz

Introduction

Confidentiality and authenticity of user files can be provided by encryption on several layers.

Application can implement encryption directly, examples are PGP or ZIP with password.

Encryption of files (inside filesystem or through independent layer like Linux eCryptfs) provides more generic solution. But some parts (like filesystem metadata) are still unencrypted. However this solution provides easy way how to store data encrypted with private key per user. (Every user can have own directory encrypted by own key.)

The lower layer is the encryption of the whole disk called Full Disk Encryption. FDE is completely transparent to the user (user does not need to choose what to encrypt because the whole disk is encrypted).

The disadvantage is that everyone who knows the password can read the whole disk, so for encryption private to user it must be combined with another layer of encryption.

The primary use for FDE is prevention to data loss in power-down mode (stolen laptop). Once the disk is unlocked, the main encryption key remains in system, usually in system RAM.

Another disadvantage of the FDE is that it cannot guarantee integrity of data. Because the encryption is fully transparent, the ciphertext and plaintext device are of the same size, there is no space to store any integrity information. This allows various attacks by direct modification of ciphertext.

The FDE works on the sector level, as the same as the block device. The atomic units for the encrypted device are sectors. Sectors are encrypted independently with symmetric cipher. Because size of symmetric cipher block is smaller than size of sector, an encryption mode have to be used.

It is crucial that the same data in different sectors must produce different ciphertext pattern. This is achieved by a per-sector tweak (IV, Initialization Vector) which is usually derived from logical sector offset (sector number).

Combination of properly used IV and encryption mode is critical for the security of the FDE system.

Figure 1. shows example of wrongly used mode for encryption (here ECB, Electronic Codebook, or in other words encryption without sector tweak and without feedback). Another problems could appear with CBC mode [CBC] and predictable IV. Then attacker can use predictability of IV and create watermarks and detect them directly from ciphertext (to prove existence of some data). Today, the most used encryption mode for disks is XTS [XTS]. This mode is constructed such way, that it can safely use predictable sector number as IV.



Figure 1. Result of statistical analysis on image encrypted in ECB mode

The FDE can be implemented on hardware layer (disks encrypted internally, encryption inside controller chipset) or in software [FDE]. Software implementation is for example Bitlocker for Windows systems, dm-crypt in Linux, FileVault on MacOS or TrueCrypt as a multiplatform solution. There are many other tools implementing this functionality [COMP].

In the following exercises we will use two software FDE tools. The first one is LUKS/dm-crypt [LUKS] which is primary solution for Linux based systems and TrueCrypt [TC].

Both tools allows a user to create a virtual encrypted block device within a partition or a regular file. This virtual disk can be formatted as if it was a physical device and mounted. TrueCrypt also supports trivial steganography and allows to create a hidden virtual volume within another volume.

References

[FDE] Disk encryption

http://en.wikipedia.org/wiki/Disk_encryption

[COMP] Comparison of disk encryption software

http://en.wikipedia.org/wiki/Comparison_of_disk_encryption_software

[CBC] Cipher Block Chaining mode

http://en.wikipedia.org/wiki/Block_cipher_mode_of_operation#Cipher-block_chaining_.28CBC.29

[XTS] IEEE P1619™/D16 Standard for Cryptographic Protection of Data on Block-Oriented Storage Devices (AES-XTS)

<http://grouper.ieee.org/groups/1619/email/pdf00086.pdf>

[LUKS] Linux Unified Key Setup / cryptsetup

<https://code.google.com/p/cryptsetup/>

[TC] TrueCrypt, Free open-source disk encryption software

<http://www.truecrypt.org/>

Exercises

During the exercises, we will use prepared Virtual Machine PV204 (Ubuntu 14.04). VM can be run in Oracle VirtualBox.

The virtual machine is preinstalled with LUKS disk encryption of the whole system, **cryptsetup** 1.6 (which allows to map TrueCrypt devices directly) and **TrueCrypt** 7.1a. For analysis of raw data you can use **okteta** hexa editor.

For experiments you can use **/dev/sdb** disk, which is 256MB virtual device unused for system.

User name: pv204

Password (also unlocking password for disk): **pv204**

Note: User pv204 is a member of group admin and can run superuser commands via sudo without need of entering password, e.g. `sudo ls /root`.

Exercise I: Hardware keylogger use

"If you let your machine out of your sight, it's no longer your machine."

In a context of disk encryption specialized hw device can be used to attack entered passphrase (BIOS password, disk encryption unlocking passphrase).

Hardware keylogger is a device which intercepts data entered through keyboard and store it internally for later analysis. The device can also transmit data using wireless network or another hidden channel. If properly masked (e.g. in keyboard controller) it is very hard to reveal it by user because device is completely transparent to BIOS and operating system (and power consumption is minimal).

USB keylogger (example: KeyDaemon module)




Figure 2. Example keylogger module

We will use simple example of commercially available USB keylogger module (2010) intended for built-in applications. Used version has only 4MB internal memory for log but extended version provides SD card for log store (enough storage for almost the whole physical life of the keyboard).

Log retrieval is activated by an activation shortcut, which switches keylogger into USB drive emulation mode and provides direct access to log storage (file).

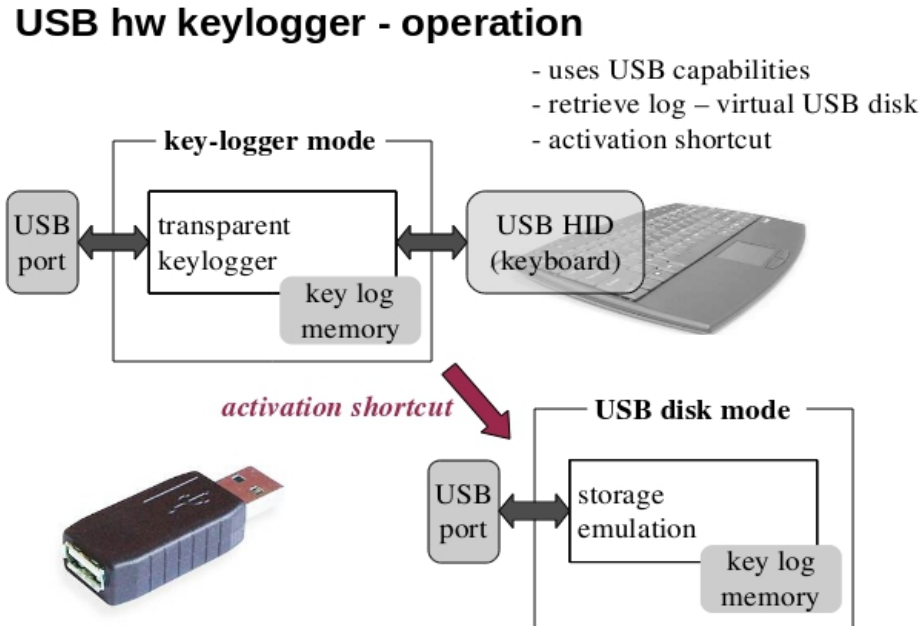


Figure 3. Log retrieval from the module

Your task

1. Connect provided USB keyboard through cable with USB keylogger module.
2. Check that system recognizes keyboard (as the same as it is connected directly).
3. Simulate some work where you need enter sensitive data, e.g. create new TrueCrypt encrypted device, boot testing virtual machine and log in or try to enter some password in browser.
4. Switch to host system file explorer and press activation shortcut :
'K' + 'B' + 'S' key together.
The keyboard should detach and you should see new virtual USB flash disk.
5. Investigate "log.txt" file on the disk and compare it with your work in step 3.
6. Please delete the "log.txt" so your colleagues will not see your "sensitive" data :-)

Exercise II: Obtaining encryption key from VM memory image

For most of the software-based disk encryption configurations the main encryption run on the main CPU, usually inside OS context (optionally with with HW acceleration like AES-NI instructions).

In this operating mode the encryption key is present in RAM during the operation. The attacker can obtain image of the physical memory and try to search for encryption key. This image can be either obtained through software (we will use VirtualBox debugger function) or through hardware (Cold boot attack, i.e. recover memory content after force reboot/short powerdown or use FireWire hw debug functions). The obtained key can be directly used for storage decryption (without password knowledge).

Next problem is to identify candidate keys in memory image. This can be done by recognizing internal OS structures in memory or by a more generic approach, like reconstructing AES keys by identifying specific round keys [COLD].

Your Task

1. Prepare Virtual Machine with some active and mounted encrypted devices and understand the storage stack inside.

You should have at least two encrypted and active disks. One is system encryption itself (LUKS/dm-crypt) and second will be TrueCrypt device created on **/dev/sdb**.

- a) Start and login to VM.
- b) Run TrueCrypt GUI (on launcher) and create new device directly on **/dev/sdb** (please use AES only for encryption algorithm). (Create volume within a partition/drive, standard volume, select device **/dev/sdb**, use AES.)
- c) mount device through TrueCrypt GUI and investigate used parameters (Volume Properties). Keep it mounted to investigate key in memory later.
- d) Investigate and understand storage stack, dump parameters about encrypted virtual devices. Use **lsblk** command to get the whole picture:

```
pv204@ubuntu32:~$ lsblk
NAME                                MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sda                                  8:0    0   8G  0 disk
├─sda1                               8:1    0 243M  0 part  /boot
├─sda2                               8:2    0    1K  0 part
├─sda5                               8:5    0  7,8G  0 part
│   └─sda5_crypt (dm-0)              252:0    0  7,8G  0 crypt
│       └─ubuntu--vg-root (dm-1)     252:1    0  6,7G  0 lvm    /
│           └─ubuntu--vg-swap_1 (dm-2) 252:2    0 1020M  0 lvm    [SWAP]
└─sdb                                8:16    0 256M  0 disk
    └─truecrypt1 (dm-3)              252:3    0 255,8M  0 dm     /media/truecrypt1
```

So **sda5** is your system encrypted partition, **sdb** is the TrueCrypt device.

Check TrueCrypt device - note the header is encrypted so you cannot get information using **blkid**. Use cryptsetup TrueCrypt extension to get more info about device on-disk header:

```
pv204@ubuntu32:~$ sudo blkid /dev/sdb
pv204@ubuntu32:~$ sudo cryptsetup tcryptDump /dev/sdb --dump-master-key -q
Enter passphrase:
TCRYPT header information for /dev/sdb
Cipher chain:  aes
Cipher mode:   xts-plain64
Payload offset: 256
MK bits:      512
MK dump:      2e 3c 6c c1 0e f8 c1 f3 ef 2d 3b 66 5d 56 aa 62
              79 6a 4e 91 c9 bf 94 6f 02 1e d5 e1 6f 7c a5 9f
              14 0a a1 15 27 4e 09 a1 73 a0 e0 6d f1 d5 4b 6e
              34 bf 4f 4e 61 98 ed 71 3d 80 0b f2 e9 85 c4 04
```

Now, try to display encrypted keys directly from kernel (superuser can display full device-mapper dm-crypt mapping table for active devices).

```
pv204@ubuntu32:~$ sudo dmsetup table --showkeys truecrypt1
0 523776 crypt aes-xts-plain64 2e3c6cc10ef8c1f3ef2d3b665d56aa62796a4e91c9bf946f021ed5e1
6f7ca59f140aa115274e09a173a0e06df1d54b6e34bf4f4e6198ed713d800bf2e985c404 256 8:16 256
```

e) Repeat the same exercise for LUKS device (**/dev/sda5** mapped to **sda5_crypt**), just use **luksDump** command. Note LUKS has visible header on-disk.

2. While the VM is still running, obtain snapshot of memory core image.

You can use **Vbox_save_memcore.bat** script which uses internal VirtualBox debugger to dump VM memory core (it is in fact ELF format dump, but it is enough for our use).

3. Analyze the memory core image with the provided **aeskeyfind** program. Note this program search for specific AES structures. (In reality you have key in memory more times but in different format, just think about screen buffer with text dumps executed in step 1.)
4. Compare possible encryption keys with information you obtained in step 1. See XTS encryption mode definition [XTS] and think why you see two separate AES keys for one device.

References

[COLD] *Lest We Remember*: Cold Boot Attacks on Encryption Keys
<https://citp.princeton.edu/research/memory/>

[XTS] IEEE P1619™/D16 Standard for Cryptographic Protection of Data on Block-Oriented Storage Devices (AES-XTS)
<http://grouper.ieee.org/groups/1619/email/pdf00086.pdf>

Exercise III: Revealing TrueCrypt hidden disk existence (CBC mode)

The goal of this exercise is to show that old CBC mode in TrueCrypt is vulnerable to watermarking attack and that such attack could be used to reveal hidden disk existence just from ciphertext analysis.

Theory

The TrueCrypt (in version before 4.1) used CBC mode with Initialization Vector calculated from the sector number and xored with secret key K_{IV} . Also there is additional whitening which is calculated using several CRC32 operations over secret key K_W xored with sector number. Final whitening value is then applied to the whole plaintext sector with plain xor function. Unfortunately, both operations are not secure (IV is still partially predictable because for consecutive sectors we know which bits will change, and the whitening is just linear transformation, CRC32 is not cryptographically secure). The figure shows simplified attack with P (plaintext block) and C (ciphertext block). For more info see [TC4].

CBC mode – hidden disk watermark

- for sector S_n divisible by 4: $S_n \text{ xor } S_{n+1} = S_{n+2} \text{ xor } S_{n+3}$
- whitening W : $W(K_w, A \text{ xor } B) = W(K_w, A) \text{ xor } W(K_w, B)$
- let's change the first plaintext blocks of sectors $n \dots n+3$
 P_n and $P_{n+2} = 0$
 P_{n+1} and $P_{n+3} = S_n \text{ xor } S_{n+1}$
- this allows to eliminate encryption and we can show that

$$C_n = E(K_E, P_n \text{ xor } IV_n) \text{ xor } W_n = E(K, K_{iv} \text{ xor } S_n) \text{ xor } W_n$$

$$C_{n+1} = E(K_E, P_{n+1} \text{ xor } IV_{n+1}) \text{ xor } W_{n+1} = \dots = E(K, K_{iv} \text{ xor } S_n) \text{ xor } W_{n+1}$$

$$C_n \text{ xor } C_{n+1} = C_{n+2} \text{ xor } C_{n+3}$$

Watermark detection
only from ciphertext!

This attack allows construct a special plaintext, which if aligned to proper sectors can propagate special pattern (watermark) into ciphertext. Because filesystem aligns start of files to sector offset, we can just use specially formatted file in hidden disk and then search for the watermark on ciphertext device. As the exercise will show, this special file can be text file, so forcing user to store such file in the hidden disk is not too complicated (imagine mail or picture in browser cache).

Your Task

1. Use the provided TrueCrypt container **tc_hidden_cbc.tc** and try to open it in TrueCrypt. The password to outer (decoy) volume is "password", password to hidden volume is "hidden".
2. Using the provided program **create_file.exe** (which implements attack above) generate special file and copy it to hidden TrueCrypt disk.
3. Dismount the TrueCrypt device and run **detect_file.exe** over the TrueCrypt container image. It should find the pattern in special file and print the offset of this pattern.
4. See the create_file and detect_file source code to better understand internal operation.

References

- [TC4] sci-crypt mailinglist, *TrueCrypt 4.0 Out*
[https://groups.google.com/forum/#!topic/sci.crypt/3DxOChZ0lrQ\[1-25-false\]](https://groups.google.com/forum/#!topic/sci.crypt/3DxOChZ0lrQ[1-25-false])

Assignment

There should be **assignment.zip** in study materials which contains **pv204_assignment.tc** (TrueCrypt virtual volume compatible with TrueCrypt 7.1a).

The volume is protected by a 9-character long password, which begins with "**pv204_XXX**" where X means digit [0-9].

Your task

1. Find the password and unlock the volume.
 2. Investigate the master encryption keys and primary header salt. What's wrong with them?
 3. Describe the findings in text file:
 - The password to the volume.
 - How did you find the password.
Attach the script or source code of your password search utility.
 - Describe which cipher chain (names of ciphers) is used for the volume.
 - Write in hexa format the volume **salt** (for the primary TrueCrypt header, from the first sector).
 - Write in hexa **master keys** for all used ciphers.
 - Shortly describe (max. 10 sentences) what is wrong with the keys (or salt).
Are the keys independent and randomly generated? If not, how do you think it was generated?
- Pack the description text file and scripts or source files to **FDE_<UCO>.zip** archive and upload it to IS.

Notes

The volume could use cipher chain (more cipher in sequence), note the most password crackers do not support this mode.

You can simply use wrapper script for TrueCrypt commandline or generate candidate password list and use dictionary search example in cryptsetup sources.

To display hexa keys you can use commands from Exercise II.

To display salt just check TrueCrypt header format [TCH] (to see the salt position) and use hexa editor, hexdump or similar tool.

References

[TCH] TrueCrypt Volume format specification
<http://www.truecrypt.org/docs/volume-format-specification>