

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

Hiding large amounts of data in virtual disk images

Steganography on virtual disks

Gašper Fele-Žorž Andrej Brodnik

University of Ljubljana, Faculty of Computer and Information Science
{polz,andy}@fri.uni-lj.si

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

- ▶ Soldiers relaying secret data¹
- ▶ Trade union negotiations¹
- ▶ Power abuse by police¹
- ▶ Robbers stealing credentials¹
- ▶ Watermarking

¹Ross Anderson, Roger Needham, and Adi Shamir (1998). “The steganographic file system”. In: International Workshop on Information Hiding. Springer, pp. 73–82.

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

- ▶ Viruses
- ▶ Sensitive data
- ▶ Illegal data

We need forensic tools!

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

- ▶ Least significant bits in images
- ▶ DCT coefficients in JPEG
- ▶ In sound files

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general


Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools



(a) Lena Cover, File Size = 44KB (b) Lena Stego, File Size = 44KB, Payload= 5019 Bytes



²Mahendra Kumar and Richard E. Newman (2010). "J3: High payload histogram neutral JPEG steganography". In: 2010 Eighth International Conference on Privacy, Security and Trust, pp. 46–53.

Hiding
large
amounts of
data in
virtual disk
images

► in ZIP files³

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

Local File Header harmless_project_documentation.doc
Local File Header boring_manual.pdf
Local File Header old_source_code.cs
Central Directory pointer to harmless_project_documentation pointer to boring_manual pointer to old_source_code



Local File Header harmless_project_documentation.doc
Local File Header boring_manual.pdf
Local File Header inofficial_backdoor_information.txt
Local File Header old_source_code.cs
Central Directory pointer to harmless_project_documentation pointer to boring_manual pointer to old_source_code

³Corina John (2006). Steganography 16 – Hiding additional files in a ZIP archive. CodeProject. last accessed May 23rd, 2017. URL: <http://www.codeproject.com/Articles/13808/Steganography-Hiding-additional-files-in-a-ZIP>.

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

► In HTML / XML Files⁴

Steganographic
filesystems

► Can you tell Cyrillic from Latin?

Hiding data
in virtual
images

Forensic
tools

⁴Sandipan Dey, Hameed Al-Qaheri, and Sugata Sanyal (2009).
“Embedding Secret Data in HTML Web Page”. In: Image Processing
and Communications Challenges. Academy Publishing House EXIT,
pp. 474–481.



Steganographic filesystems

Hiding
large
amounts of
data in
virtual disk
images

- ▶ StegFS - on Linux⁵
- ▶ StegFS - on Windows⁶

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

⁵Andrew D. McDonald and Markus G. Kuhn (2000). “StegFS: A Steganographic File System for Linux”. In: Information Hiding: Third International Workshop, IH’99, Dresden, Germany, September 29 - October 1, 1999 Proceedings. Ed. by Andreas Pfitzmann. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 463–477. ISBN: 978-3-540-46514-0. DOI: 10.1007/10719724_32. URL: http://dx.doi.org/10.1007/10719724_32.

⁶Hungseok Pang, Kian-Lee Tan, and Xiaolin Zhou (2003). “StegFS: a steganographic file system”. In: Proceedings of the 19th International Conference on Data Engineering. IEEE, pp. 657–667. DOI: 10.1109/ICDE.2003.1260829.

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

- ▶ Raw
- ▶ Pre-allocated
- ▶ Sparse (QCOW, VDI, VMDK, VHD)

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

- ▶ Header
- ▶ Cluster table
- ▶ Clusters

3 ways to hide data

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

- ▶ In headers
- ▶ In backing files
- ▶ Between clusters

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

Disk DescriptorFile

version=1

CID=bee9efa

parentCID=ffffff

createType="monolithicSparse"

Extent description

RW 6291936 SPARSE "test1.vmdk"

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

Parent




Figure: Using differencing virtual images to hide data. The dark gray sector is hidden from the user. The light gray sectors are not present in the Child disk image - they are only present in the Parent

Hiding
large
amounts of
data in
virtual disk
images


Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

Child



Parent




Figure: Using differencing virtual images to hide data. The dark gray sector is hidden from the user. The light gray sectors are not present in the Child disk image - they are only present in the Parent

Hiding
large
amounts of
data in
virtual disk
images


Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

Child



Parent




Figure: Using differencing virtual images to hide data. The dark gray sector is hidden from the user. The light gray sectors are not present in the Child disk image - they are only present in the Parent

Hiding
large
amounts of
data in
virtual disk
images


Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

Child



Parent




Figure: Using differencing virtual images to hide data. The dark gray sector is hidden from the user. The light gray sectors are not present in the Child disk image - they are only present in the Parent

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

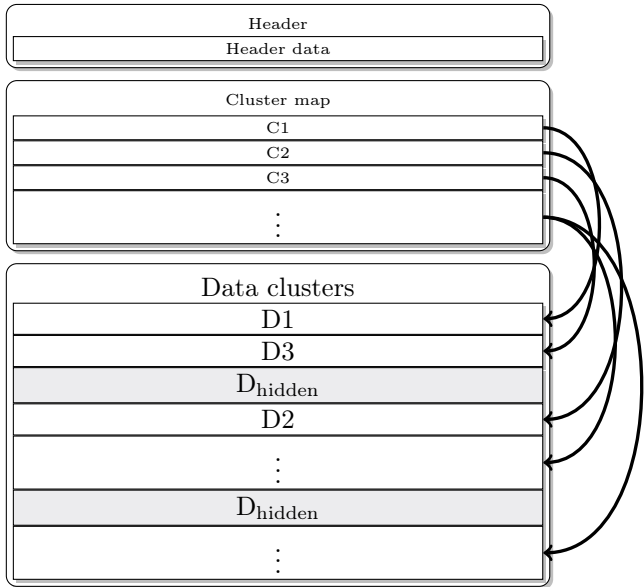


Figure: Hiding data between clusters

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

Detecting/removing hidden data:

- ▶ Convert to raw and back
- ▶ Compare file sizes
- ▶ Initial version:
http://github.com/polz113/virtual_disk_injector



Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

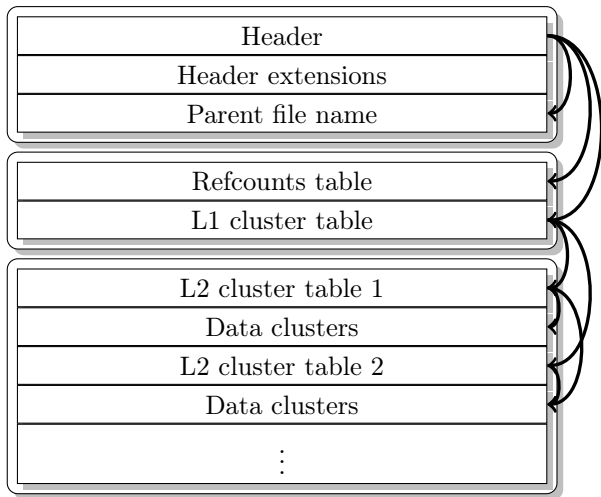


Figure: The structure of a QCOW2 file. Arrows represent pointers from one section to another.

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

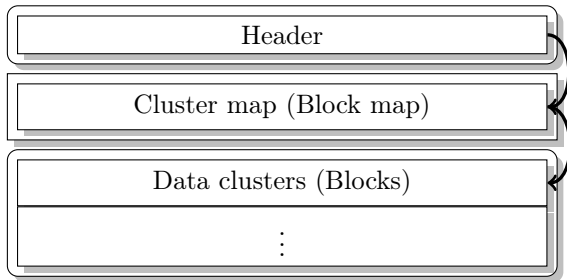


Figure: The structure of a VDI file. Arrows represent pointers from one section to another.

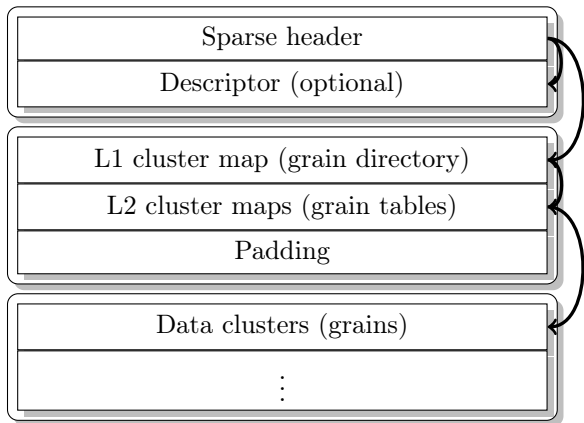


Figure: The structure of a VMDK file. Arrows represent pointers from one section to another. Starts of data clusters are aligned to cluster size.

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

Hiding
large
amounts of
data in
virtual disk
images

Steganography
in general

Steganographic
filesystems

Hiding data
in virtual
images

Forensic
tools

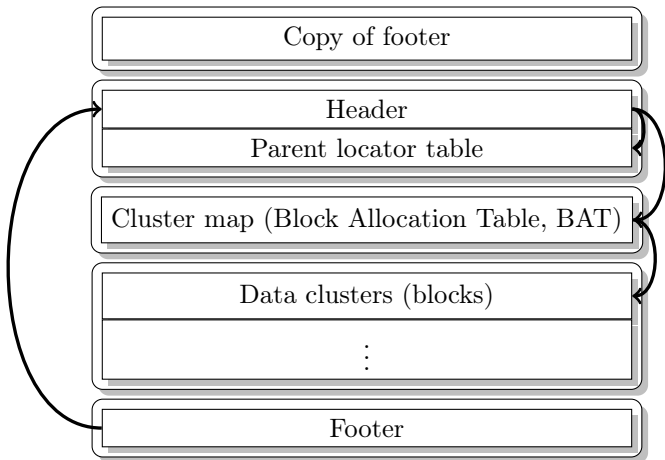


Figure: The structure of a VHD file. Arrows represent pointers from one section to another.