



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Steganography

While classical cryptography is about concealing the content of messages, steganography is about concealing their existence.

Classical steganography concerns itself with ways of embedding a secret message (which might be a copyright mark, or a covert communication, or a serial number) in a cover message (such as a video film, an audio recording, or computer code). The embedding is typically parametrised by a key; without knowledge of this key (or a related one) it is difficult for a third party to detect or remove the embedded material. Once the cover object has material embedded in it, it is called a stego object. Thus, for example, we might embed a mark in a covertext giving a stegotext; or embed a text in a cover image giving a stego-image; and so on. (This terminology was agreed at the First International Workshop on Information Hiding).ⁱ

History

Throughout history, people have hidden information by a multitude of methods and variations. For example, ancient Greeks wrote text on wax-covered tablets. To pass a hidden message, a person would scrape wax off a tablet, write a message on the underlying wood and again cover the tablet with wax to make it appear blank and unused. Another ingenious method was to shave the head of a messenger and tattoo a message or image on the messenger's head. After the hair grew back, the message would be undetected until the head was shaved again.

Invisible inks offered a common form of invisible writing. Early in World War II, steganographic technology consisted almost exclusively of these inks. With invisible ink, a seemingly innocent letter could contain a very different message written between the lines.

Documents themselves can hide information: document text can conceal a hidden message through the use of null ciphers (unencrypted messages), which camouflage the real message in an innocent-sounding missive. Open coded messages, which are plain text passages, "sound" innocent because they purport to be about ordinary occurrences. Because many open-coded messages do not seem to be cause for suspicion, and therefore "sound" normal and innocent, the suspect communications can be detected by mail filters while "innocent" messages are allowed to flow through. For example, the following null-cipher message was actually sent by a German spy in WWII:

Apparently neutral's protest is thoroughly discounted and ignored. Isman hard hit. Blockade issue affects pretext for embargo on by-products, ejecting suets and vegetable oils.

Decoding this message by extracting the second letter in each word reveals the following, hidden message:

Pershing sails from NY June 1.

Document layout may also reveal information. Documents can be marked and identified by modulating the position of lines and words.

Message detection improved with the development of new technologies that could pass more information and be even less conspicuous. The Germans developed microdot technology, which FBI director J. Edgar Hoover referred to as “the enemy’s masterpiece of espionage.” Microdots are photographs the size of a printed period having the clarity of standard-sized typewritten pages, which permits the transmission of large amounts of data, including drawing and photographs.

With every discovery of a message hidden with an existing application, a new steganographic application is being devised. Old methods are given new twists. While drawings have often been used to conceal or reveal information, computer technology has, in fact, sparked a revolution in such methods for hiding information.¹¹

Steganography Today

In the present day, digital images, (as well as audio and video files) offer a rich environment for hiding virtually unlimited types of data. To illustrate, the following example shows two images whose differences are undetectable with the naked eye. The image on the right, however, contains a secret message in Microsoft® Word format that is 88KB in size.



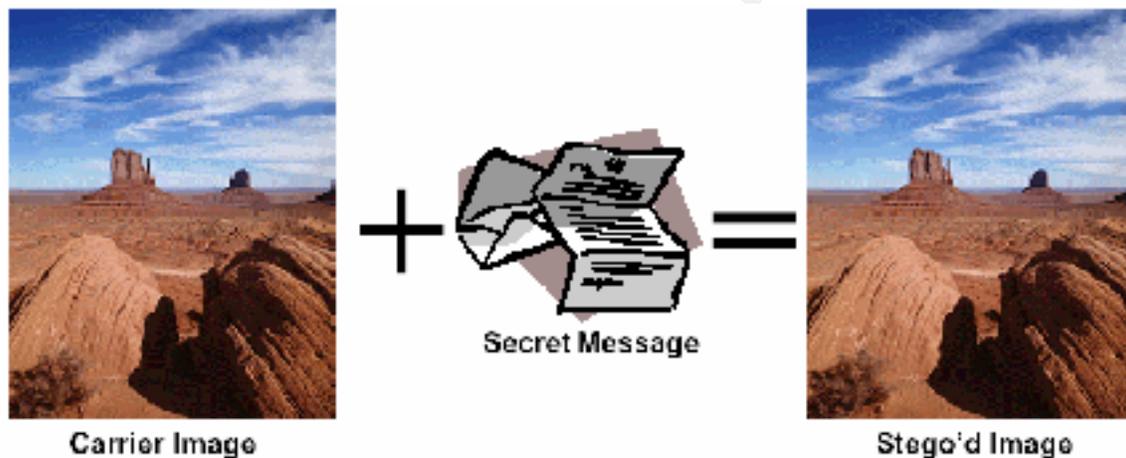
Original



Original + Secret Message

The proliferation of steganography software is widespread. During a recent research effort funded by the U.S. Air Force Research Laboratory, WetStone Technologies examined over 50 programs that hide messages in image files, audio and video files, and in text documents. Many of these programs are freely downloadable from the Internet, while some charge a nominal fee of between \$49 and \$89. Research indicates that a minimum of 1 million copies of steganography software have been downloaded or purchased over the Internet during the last 18 months. It is likely that this is a conservative estimate.

There are many techniques for hiding secret messages in images. The common denominator among them is that they combine a carrier file with a secret message in order to produce a new image containing a hidden message that is indiscernible to the eye. The methodologies vary widely from simple least significant bit modification [See notes for information on the least significant bit] to quite sophisticated JPEG coefficient transforms [See notes section for information on JPEG coefficient transforms].ⁱⁱⁱ



STEGANALYSIS

Hiding information within electronic media requires alterations of the media properties that may introduce some form of degradation or unusual characteristics. These characteristics may act as signatures that broadcast the existence of the embedded message, thus defeating the purpose of steganography. Attacks and analysis on hidden information may take several forms: detecting, extracting, and disabling or destroying hidden information. An attacker may also embed counter-information over the existing hidden information. Due to space limitations we will look at two methods: detecting messages or their transmission and disabling embedded information. These approaches (attacks) vary depending upon the methods used to embed the information into the cover media. Our goal is not to advocate the removal or disabling of valid hidden information such as copyrights, but to point out approaches that are vulnerable and may be exploited to investigate illicit hidden information. Some amount of distortion and degradation may occur to carriers of hidden messages even though such distortions cannot be detected easily by the human perceptible system. This distortion may be anomalous to the

"normal" carrier that when discovered may point to the existence of hidden information. Steganography tools vary in their approaches for hiding information. Without knowing which tool is used and which, if any, stegokey is used; detecting the hidden information may become quite complex. However, some of the steganographic approaches have characteristics that act as signatures for the method or tool used.

DETECTING HIDDEN INFORMATION

Unusual patterns stand out and expose the possibility of hidden information. In text, small shifts in word and line spacing may be somewhat difficult to detect by the casual observer. However, appended spaces and "invisible" characters can be easily revealed by opening the file with a common word processor. The text may look "normal" if typed out on the screen, but if the file is opened in a word processor, the spaces, tabs, and other characters distort the text's presentation.

Images too may display distortions from hidden information. Selecting the proper combination of steganography tools and carriers is key to successful information hiding. Some images may become grossly degraded with even small amounts of embedded information. This "visible noise" will give away the existence of hidden information. The same is true with audio. Echoes and shadow signals reduce the chance of audible noise, but they can be detected with little processing.

Only after evaluating many original images and stegoimages as to color composition, luminance, and pixel relationships do anomalies point to characteristics that are not "normal" in other images. Patterns become visible when evaluating many images used for applying steganography. Such patterns are unusual sorting of color palettes, relationships between colors in color indexes, exaggerated "noise".

An approach used to identify such patterns is to compare the original cover-images with the stego-images and note visible differences (known-cover attack). Minute changes are readily noticeable when comparing the cover and stego-images. In making these comparisons with numerous images, patterns begin to emerge as possible signatures of steganography software. Some of these signatures may be exploited automatically to identify the existence of hidden messages and even the tools used in embedding the messages. With this knowledge base, if the cover images are not available for comparison, the derived known signatures are enough to imply the existence of a message and identify the tool used to embed the message. However, in some cases recurring, predictable patterns are not readily apparent even if distortion between the cover and stego-images is noticeable.

A number of disk analysis utilities are available that can report and filter on hidden information in unused clusters or partitions of storage devices. A steganographic file system may also be vulnerable to detection through analysis of the systems partition information.

Filters can also be applied to capture TCP/IP packets that contain hidden or invalid information in the packet headers. Internet firewalls are becoming more sophisticated and allow for much customization. Just as filters can be set to determine if packets originate from within the firewall's domain and the validity of the SYN and ACK bits, so too can the filters be configured to catch packets that have information in supposed unused or reserved space.

DISABLING STEGANOGRAPHY

Detecting the existence of hidden information defeats the steganography's goal of imperceptibility. Methods exist that produce results which are far more difficult to detect without the original image for comparison. At times the existence of hidden information may be known so detecting it is not always necessary. Disabling and rendering it useless seems to be the next best alternative. With each method of hiding information there is a trade off between the size of the payload (amount of hidden information) that can be embedded and the survivability or robustness of that information to manipulation.

The distortions in text noted by appended spaces and "invisible" characters can be easily revealed by opening the file with a word processor. Extra spaces and characters can be quickly stripped from text documents.

The disabling or removal of hidden information in images comes down to image processing techniques. For LSB methods of inserting data, simply using a lossy compression technique [See notes section for information on lossy compression], such as JPEG, is enough to render the embedded message useless. Images compressed with such a method are still pleasing to the human eye but no longer contain the hidden information

Tools exist to test the robustness of information hiding techniques in images. These tools automate image-processing techniques such as warping, cropping, rotating, and blurring. Such tools and techniques should be used by those considering making the investment of watermarking to provide a sense of security of copyright and licensing just as password cracking tools are used by system administrators to test the strength of user and system passwords. If the password fails, the administrator should notify the password owner that the password is not secure.

Hidden information may also be overwritten. If information is added to some media such that the added information cannot be detected, then there exists some amount of additional information that may be added or removed within the same threshold which will overwrite or remove the embedded covert information

Audio and video are vulnerable to the same methods of disabling as with images. Manipulation of the signals will alter embedded signals in the noise level (LSB), which may be enough to overwrite or destroy the embedded message. Filters can be used in an attempt to cancel out echoes or subtle signals but becomes this may not be as successful as expected.

Caution must be used in hiding information in unused space in files or file systems. File headers and reserved spaces are common places to look for “out of place” information. In file systems, unless the steganographic areas are in some way protected (as in a partition), the operating system may freely overwrite the hidden data since the clusters are thought to be free. This is a particular annoyance of operating systems that do a lot of caching and creating of temporary files. Utilities are also available which “clean” or wipe unused storage areas. In wiping, clusters are overwritten several times to ensure any data has been removed. Even in this extreme case, utilities exist that may recover portions of the overwritten information.

As with unused or reserved space in file headers, TCP/IP packet headers can also be reviewed easily. Just as firewall filters are set to test the validity of the source and destination IP addresses, the SYN and ACK bits, so too can the filters be configured to catch packets that have information in supposed unused or reserved space. If IP addresses are altered or spoofed to pass covert information, a reverse lookup in a domain name service (DNS) can verify the address. If the IP address is false, the packet can be terminated. Using this technique to hide information is risky as TCP/IP headers may get overwritten in the routing process. Reserved bits can be overwritten and passed along without impacting the routing of the packet.^{iv}

Conclusion

While steganography is in no way a new science, it is emerging with the prevalence of better technology as one of the most interesting and potentially helpful technologies today. When properly combined with cryptography, steganography can offer a secure means of communication, digital copywriting, protection of intellectual property, development of secure covert communications channels as well as many anti-tampering issues.

On the other hand steganography’s potential to hide information also plays well into the hands of criminals. While it is very difficult to detect if an image or audio file has been altered to contain a stego bit of information; it is not impossible, and only through security professionals learning more about this very exciting field can we better understand how to utilize steganography in beneficial ways and to thwart its use for the criminal elements.

To get a good hands on introduction or to have some fun with steganography, try a free product called S-Tools. Also check this page for other tools you can use to experiment with steganography. <http://members.tripod.com/steganography/stego/software.html>

Notes

The Least Significant Bit is the lowest-order bit. For example, the least significant bit of 01000111 is a 1 (that's the 1 on the very right). The most significant bit is a 0 (yep, that's the 0 on the very left).^v

JPEG coefficient transforms refers to the discrete cosine transform (DCT),^{vi} which helps to separate the image into parts (or spectral sub-bands) of differing importance (with respect to the image's visual quality)).

Lossless and lossy compression are terms that describe whether or not, in the compression of a file, all original data can be recovered when the file is uncompressed. With lossless compression, every single bit of data that was originally in the file remains after the file is uncompressed. All of the information is completely restored. This is generally the technique of choice for text or spreadsheet files, where losing words or financial data could pose a problem. The Graphics Interchange File (GIF) is an image format used on the Web that provides lossless compression.

On the other hand, lossy compression reduces a file by permanently eliminating certain information, especially redundant information. When the file is uncompressed, only a part of the original information is still there (although the user may not notice it). Lossy compression is generally used for video and sound, where a certain amount of information loss will not be detected by most users. The JPEG image file, commonly used for photographs and other complex still images on the Web, is an image that has lossy compression. Using JPEG compression, the creator can decide how much loss to introduce and make a trade-off between file size and image quality.^{vii}

ⁱ Taken from On The Limits of Steganography by Ross J. Anderson and Fabien A.P. Petitcolas – IEEE Journal of Selected Areas in Communications, 16(4):474-481, May 1998. Special Issue on Copyright & Privacy Protection. ISSN 0733-8716 –

ⁱⁱ Taken from Exploring Steganography: Seeing the Unseen by Neil F. Johnson and Sushil Jajodia George Mason University - <http://www.jjtc.com/pub/r2026.pdf>

ⁱⁱⁱ Taken from What you can't see can Hurt You . . . The Dangers of Steganography by Chet Hosmer, President & CEO WetStone Technologies, Inc. – <http://www.wetstonetech.com/stego-paper.pdf>

^{iv} Taken from Steganalysis: The Investigation of Hidden Information by Neil F. Johnson and Sushil Jajodia – Center for Secure Information Systems, George Mason University, MS:4A4, Fairfax, Virginia 22030-4444 - <http://www.simovits.com/archive/it98jjgmu.pdf>

^v http://www.oreilly.com/reference/dictionary/terms/L/Least_Significant_Bit_or_Byte.htm

^{vi} <http://www.ece.purdue.edu/~ace/jpeg-tut/jpgdct1.html>

^{vii} http://whatis.techtarget.com/definition/0..sid9_gci214453.00.html