




Defense Cyber Crime Institute



# Test Report for

## StegAlyzerSS V2.0

October 2006

  
\_\_\_\_\_  
Mark Hirsh  
System Engineer

10/23/06  
\_\_\_\_\_  
Date

  
\_\_\_\_\_  
Edmond Kong  
Director, RDT&E

23 OCT 06  
\_\_\_\_\_  
Date

Defense Cyber Crime Institute

## TABLE OF CONTENTS

<b>EXECUTIVE SUMMARY .....</b>	<b>III</b>
<b>1. SCOPE .....</b>	<b>1</b>
1.1 Identification.....	1
1.2 StegAlyzerSS Features and Capabilities .....	1
1.3 Test Approach.....	1
<b>2. TEST OF SIGNATURE SEARCH FEATURE.....</b>	<b>4</b>
2.1 Test Description .....	4
2.2 Test Data .....	4
<b>3. SUMMARY OF FINDINGS FOR STEGALYZERSS V2.0.....</b>	<b>5</b>
<b>4. VENDOR COMMENTS.....</b>	<b>6</b>
<b>APPENDIX. EFFECTIVE/INEFFECTIVE STEGANALYSIS ALGORITHMS.....</b>	<b>7</b>

## **Defense Cyber Crime Institute**

### **EXECUTIVE SUMMARY**

This report describes the tests and procedures that were used to evaluate Steganography Analyzer Signature Scanner (StegAlyzerSS), v2.0. StegAlyzerSS, a steganalysis tool that uses signature analysis to search for hidden steganographic data, is a product of the Backbone Security.com Steganography Analysis and Research Center (SARC).

DCCI testing found that StegAlyzerSS can perform steganalysis on suspect media that is directly attached to the examination workstation and on dd and Encase images of suspect media.

The Signature Search option was found to be effective in identifying image files that contain hidden steganographic data. In the controlled test environment, 818 of the 821 image files identified as suspicious did actually contain hidden steganographic data, producing a degree of confidence of 99.6%. The algorithm is also useful in that it was able to identify image files that had been hidden with fourteen different steganography programs, and it was able to identify approximately 77% of the stego files that were created using those programs. In several instances, it was actually able to correctly identify the steganography program that was used to hide the data, and in some cases the hidden steganographic data could be extracted.

## **1. SCOPE**

### **1.1 IDENTIFICATION**

This report describes the tests and procedures that were used to evaluate Steganography Analyzer Signature Scanner (StegAlyzerSS), v2.0. StegAlyzerSS, a steganalysis tool that searches for hidden steganographic data, is a product of the Backbone Security.com Steganography Analysis and Research Center (SARC).

The Defense Cyber Crime Institute (DCCI) developed this test report. The intent of the testing is to determine whether StegAlyzerSS provides the law enforcement and forensic communities with an effective means of identifying files that contain hidden steganographic data.

### **1.2 STEGALYZERSS FEATURES AND CAPABILITIES**

StegAlyzerSS searches for hidden steganographic data using signature analysis, an approach that evaluates files to determine whether known signatures of specific steganography programs can be identified. The signatures that StegAlyzerSS looks for are those that Backbone Security personnel have discovered through extensive research on the steganography applications.

### **1.3 TEST APPROACH**

StegAlyzerSS's ability to identify suspicious files was evaluated using the DCCI stego library, which contains over 4000 clean files (files that do not contain hidden data) and approximately 3500 stego files (files that do contain hidden data). DCCI uses this stego library to test and evaluate steganalysis programs because it supplies a known, controlled environment. The accuracy of the output produced by the steganalysis program is easy to verify because it is known which files are clean, which contain hidden data, and, for files containing hidden data, what steganography program was used to hide the data.

The stego files were created using 51 different steganography programs (i.e., data hiding algorithms). The library is made up of audio (WAV and MP3) and image files (BMP, GIF, JPEG, TIFF, PNG, and PCX) and it includes both grayscale and color images. Each of the stego files was created using one of the clean files as a carrier file. As a result, a correspondence can be created by mapping every stego file back to the particular clean file that was used as the carrier. Clean files that were not used as carrier files for stego data hiding can also be readily identified.

DCCI tested the signature search feature using the color and grayscale image files contained in the DCCI stego library. When testing the Signature Search feature against the DCCI image files in the stego library, there are five possible outcomes that can be obtained:

## Defense Cyber Crime Institute

- A. A file can be recognized as clean when it contains no hidden data, and identified as suspicious when it does contain hidden data. (This outcome is considered **ideal** because it demonstrates that the steganalysis program is doing exactly what it is designed to do.)
- B. A clean file can be identified as suspicious even though it contains no hidden data, and identified as suspicious when it actually does contain hidden data. (This outcome is considered **ineffective** because it does not demonstrate that the steganalysis program actually identified the existence of hidden data.)
- C. A clean file can be identified as suspicious even though it contains no hidden data, and not identified as suspicious when it actually does contain hidden data. (This outcome is considered **unwanted** because it is exactly the opposite of what the steganalysis program is designed to do.)
- D. A clean file that does not correspond to any stego file can be identified as suspicious. (This outcome is considered **undesirable** because it provides misleading indications.)
- E. A clean file can be recognized as clean when it contains no hidden data, and not identified as suspicious when it actually does contain hidden data – for those clean files used as carrier files. (This outcome is considered **neutral** because it simply demonstrates that the steganalysis program is not yet able to detect certain types of steganographic algorithms.)

DCCI uses two computations to measure the effectiveness of a steganalysis program for law enforcement and forensic use: degree of confidence (DOC) and measure of usefulness (MOU). The DOC measures the percentage of the files deemed suspicious that fall into outcome A (**ideal**). It is computed by dividing the number of files appearing in outcome A by the sum of the number of files appearing in outcomes A, B, C, and D. DOC measures how confident one can be that a file identified as suspicious actually contains hidden steganographic data. Low DOC values imply the steganalysis algorithm produces a large number of false positives. In other words, many of the files that are identified as suspicious do not actually contain hidden data.

MOU measures how well a steganalysis algorithm is in identifying particular steganography algorithms. It is computed by dividing the number of files appearing in outcome A by the total number of files appearing in the stego library that are hidden using the steganography algorithms represented by the list of files in outcome A. MOU provides a way to measure how restrictive a steganalysis program is. In other words, for the steganography algorithms that the steganalysis algorithm is able to detect, it measures the actual percentage of files that are identified as suspicious.

## **Defense Cyber Crime Institute**

To be considered effective for law enforcement (LE) and forensic use, DCCI believes it is necessary to achieve a DOC of at least 85% **AND** an MOU of at least 50%.

Examples of what constitutes effective and ineffective steganalysis algorithms can be found in the Appendix.

## 2. TEST OF SIGNATURE SEARCH FEATURE

### 2.1 TEST DESCRIPTION

This test was repeated three times. In the first test the DCCI stego library was stored on a hard drive directly attached to the examination workstation. In the second test, a dd image of the hard drive containing the DCCI stego library was used as the source material, and in the third test an Encase image of the hard drive containing the stego library was used.

### 2.2 TEST DATA

Test case ID	SS-01
Test objective	Determine whether the StegAlyzerSS Signature Search option is able to distinguish between files that contain hidden steganographic data and those that do not.
Expected Results	The DOC for Signature Search will measure at least 85% <b>AND</b> the MOU will measure at least 50%.
Test Results	<b>Expected results were obtained.</b>
Test Procedure	StegAlyzerSS was initially run against all of the clean image files contained in the DCCI stego library and after the operation completed, StegAlyzerSS was run against all of the stego image files contained in the DCCI stego library. All files identified as suspicious were categorized as <b>ideal</b> , <b>ineffective</b> , <b>unwanted</b> , or <b>undesirable</b> .
Measure	A DOC and an MOU were computed to measure the effectiveness of the StegAlyzerSS Signature Search file identification process.
Actual Results	<p>1. 818 of the 821 files identified as suspicious fell into Category A, <b>ideal</b>; the other three fell into the category <b>undesirable</b>. As a result, the DOC computed for Signature Search was 99.6%, indicating that there is a high degree of confidence that any file identified as suspicious, when using this option, will actually contain hidden steganography data.</p> <p>2. Fourteen different steganography algorithms were represented by the 818 stego files identified as suspicious. Since a total of 1066 stego files have been created using these fourteen algorithms, the MOU for Signature Search was approximately 77%, suggesting that the Signature Search operation is effective at recognizing steganography signatures.</p> <p>3. The results were identical for all three test runs. That is, it did not matter whether the stego library resided on a hard drive that was directly attached to the examination workstation or the library was stored on the hard drive within a dd or Encase image.</p>
Anomalies	None.

### **3. SUMMARY OF FINDINGS FOR STEGALYZERSS V2.0**

DCCI testing found that StegAlyzerSS can perform steganalysis on suspect media that is directly attached to the examination workstation and on dd and Encase images of suspect media.

The Signature Search option was found to be effective in identifying image files that contain hidden steganographic data. In the controlled test environment, 818 of the 821 image files identified as suspicious did actually contain hidden steganographic data, producing a degree of confidence of 99.6%. The algorithm is also useful in that it was able to identify image files that had been hidden with fourteen different steganography programs, and it was able to identify approximately 77% of the stego files that were created using those programs. In several instances, it was actually able to correctly identify the steganography program that was used to hide the data, and in some cases the hidden steganographic data could be extracted.

#### 4. VENDOR COMMENTS

This test report suggests that StegAlyzerSS only searches for hidden steganographic data in image files. However, because the appending technique is used by some steganography applications to hide information by appending it to the end of any type of file, StegAlyzerSS actually scans every file on the disk or image, or subset thereof, selected by the user. DCCI only tested image files because their stego library does not include any other file types with appended data.

In addition to EnCase and raw image formats, StegAlyzerSS can also mount and scan images in SMART format for the users who use Linux platforms in their examinations.

While we are very pleased with the 99.6% Degree of Confidence, we believe it is important to describe why three files were identified as suspicious when, in fact, they did not contain any hidden steganographic data and how we plan to resolve that issue.

Some of the 22 signatures of steganography applications we have discovered to date are quite long. However, there are a few that are quite small. The three files that were identified as suspicious did not contain any hidden steganography data but did contain the signature of one of the applications that have a small signature which resulted in the three false positives.

The signature for this program is only two bytes long. Thus, the probability of those two bytes occurring in any given file is much higher than the probability of longer signatures. Although we do not have statistical evidence to prove that assertion, we believe it should be intuitively obvious that shorter signatures will yield a larger number of false positives (FPs) than longer signatures.

For a different application with a two byte signature, we were able to find additional information in the file unique to the application. By correlating the signature with additional information, it is possible to reduce the occurrence of FPs, possibly to zero.

Therefore, we intend to conduct additional analysis on this particular steganography program to determine if there is additional information in the stego file than can be correlated with the two byte signature to more accurately identify information hidden with this program and keep the number of FPs to an absolute minimum, and hopefully zero.

## APPENDIX. EFFECTIVE/INEFFECTIVE STEGANALYSIS ALGORITHMS

As stated in Section 1.4, there are five possible outcomes that can be obtained when a steganalysis algorithm is run against the 4000 clean and 3500 stego files that make up the DCCI stego library. These outcomes are:

- A. A file can be recognized as clean when it contains no hidden data, and identified as suspicious when it does contain hidden data. (This outcome is considered **ideal** because it demonstrates that the steganalysis program is doing exactly what it is designed to do.)
- B. A clean file can be identified as suspicious even though it contains no hidden data, and identified as suspicious when it actually does contain hidden data. (This outcome is considered **ineffective** because it does not demonstrate that the steganalysis program actually identified the existence of hidden data.)
- C. A clean file can be identified as suspicious even though it contains no hidden data, and not identified as suspicious when it actually does contain hidden data. (This outcome is considered **unwanted** because it is exactly the opposite of what the steganalysis program is designed to do.)
- D. A clean file that does not correspond to any stego file can be identified as suspicious. (This outcome is considered **undesirable** because it provides misleading indications.)
- E. A clean file can be recognized as clean when it contains no hidden data, and not identified as suspicious when it actually does contain hidden data – for those files used as carrier files. (This outcome is considered **neutral** because it simply demonstrates that the steganalysis program is not yet able to detect certain types of steganographic algorithms.)

DCCI uses two computations to measure the effectiveness of a steganalysis program for law enforcement and forensic use: degree of confidence (DOC) and measure of usefulness (MOU). The DOC measures the percentage of the files deemed suspicious that fall into outcome A (**ideal**). It is computed by dividing the number of files appearing in outcome A by the sum of the number of files appearing in outcomes A, B, C, and D. DOC measures how confident one can be that a file identified as suspicious actually contains hidden steganographic data. Low DOC values imply the steganalysis algorithm produces a large number of false positives. In other words, many of the files that are identified as suspicious do not actually contain hidden data.

MOU measures how well a steganalysis algorithm is in identifying particular steganography algorithms. It is computed by dividing the number of files appearing in outcome A by the total number of files appearing in the stego library that are hidden using the steganography algorithms represented by the list of files in outcome A. MOU provides a way to measure how restrictive a steganalysis algorithm is. In other words, for the steganography algorithms that the steganalysis program is able to detect, it measures the actual percentage of files that are identified as suspicious.

## Defense Cyber Crime Institute

To be considered effective for law enforcement and forensic use, DCCI believes it is necessary to achieve a DOC of at least 85% **AND** an MOU of at least 50%. The following examples describe how DCCI evaluates steganalysis algorithms to determine whether they are effective for law enforcement and forensic use.

1. **Description of Steganalysis Algorithm's Results:** This example evaluates an algorithm that identifies 1200 suspicious files when it is run against all files contained in the DCCI stego library, and all of the suspicious files fall into Category A, **ideal**. Furthermore, the suspicious files that are identified are associated with 20 different steganography programs. The total number of stego files in the stego library that were created using these 20 programs is 1600.

**Evaluation:** The DOC for this algorithm is 100% since no results appear in Categories B, C, or D ( $1200/1200+0+0+0$ ). The MOU is 75% ( $1200/1600$ ). Therefore, the steganalysis algorithm is considered effective for law enforcement and forensic use. The reason it is useful is that it shows a strong ability to distinguish between files that contain hidden steganographic data and files that do not. In addition, it is able to identify a significant number of the files that are hidden with the particular algorithms it is able to identify.

2. **Description of Steganalysis Algorithm's Results:** This example evaluates an algorithm that identifies all files in the DCCI stego library as suspicious. In this case, all suspicious files fall into Categories B and D, **ineffective** and **undesirable**. All suspicious stego files fall into category B because all of the stego files were created using as a carrier file one of the clean files and all of the clean files were identified as suspicious.

**Evaluation:** The DOC for this algorithm is 0% since no results appear in Category A ( $0/0+3500+0+4000$ ). The MOU is 0% (or more accurately undefined –  $0/0$ ) since no files fall into category A. Therefore, the algorithm is **not** considered effective for law enforcement and forensic use. The reason it is not useful is that it does not effectively distinguish between files that contain hidden steganographic data and files that do not. It is important to note that this algorithm identified more of the actual stego files as suspicious than the algorithm discussed in example 1 above, but the algorithm identified in example 1 is considered effective and this algorithm is not. **Effectiveness is based upon how well the steganalysis algorithm distinguishes between files that contain hidden steganographic data and files that do not; it is not based upon the total number of suspicious files identified.**

3. **Description of Steganalysis Algorithm's Results:** This example evaluates an algorithm that identifies 1 suspicious file when it is run against all files contained in the DCCI stego library, and the suspicious file falls into Category A, **ideal**. The total number of stego files in the stego library that were created using the steganography program that created the file identified as suspicious is 20.

## Defense Cyber Crime Institute

**Evaluation:** The DOC for this algorithm is 100% since no results appear in Categories B, C, or D (1/1+0+0+0). The MOU is 5% (1/20). Therefore, the algorithm is **not** considered useful. The reason it is not useful is that it is too restrictive. That is, even though it appears to only identify files that contain hidden steganographic data as suspicious, it is not a useful algorithm because too many files containing hidden steganographic data go undetected.

4. **Description of Steganalysis Algorithm's Results:** This example evaluates an algorithm that identifies 20 suspicious files when it is run against all files contained in the DCCI stego library, and all of the suspicious files fall into Category A, **ideal**. Furthermore, the suspicious files that are identified are associated with only one steganography program. The total number of stego files in the stego library that were created using the steganography program that created the files identified as suspicious is 20.

**Evaluation:** The DOC for this algorithm is 100% since no results appear in Categories B, C, or D (20/20+0+0+0). The MOU is also 100% (20/20). Therefore, the algorithm is considered effective. The reason it is effective is that it appears to be very good at identifying steganographic files that were created using a particular steganography program. When using this steganalysis program law enforcement and forensic examiners would have very strong indications whether a particular steganography program was used to hide data on media that had been seized. Note that this algorithm identifies far fewer actual stego files than the algorithm discussed in example 2, but this algorithm is considered effective and the one in example 2 is not. To reiterate, **effectiveness is based upon how well the steganalysis algorithm distinguishes between files that contain hidden steganographic data and files that do not; it is not based upon the total number of suspicious files identified.**