Counter-Forensic Tools: Analysis and Data Recovery

Matthew Geiger
PDT Forensics Team, CERT
Pittsburgh, PA
mgeiger@cert.org

Abstract

Among the challenges faced by forensic analysts are a range of commercial 'disk scrubbers', software packages designed to irretrievably erase files and records of computer activity. These counter-forensic tools have been used to eliminate evidence in criminal and civil legal proceedings and represent an area of continuing concern for forensic investigators.

This paper details the analysis of 13 commercial counter-forensic tools, examining operational shortfalls that can permit the recovery of significant evidentiary data. The research also isolates filesystem fingerprints generated when these tools are used, which can identify the tool, demonstrate its actual use and, in many cases, provide insight into the extent and time of its use.

One result is an indexed resource for forensic analysts, covering 19 tools and tool versions, that can help identify traces of disk-scrubbing activity and guide the search for residual data. In addition, a new forensic utility, named Aperio, is presented. It employs a signature library to automate the hunt for traces of counter-forensic tool use. Aperio can search filesystems presented as images or devices, and provides a detailed audit report of its findings. Together these resources may assist in establishing the usage of counter-forensic tools where such activity has legal implications.

Introduction

Modern computer operating systems and software applications generate copious amounts of data about their users' activity. These records, as well as user-created files, represent valuable sources of evidence and are increasingly the focus of investigation and legal discovery.

At the same time, users have grown more aware that "deleting" files does not mean obliterating the information they contain. Their awareness has fueled a market for counter-forensic software, which vendors promote as guarding users' privacy and/or protecting them from the consequences of their activity on the computer.

More than 25 such counter-forensic tools can be readily located via popular Internet search engines and referral-driven Web sites, such as http://www.privacy-software-review.com.

These commercial tools claim to expunge all traces of information about specific computer usage, including documents and other files created, records of websites visited, images viewed and files downloaded. To do this, counter-forensic tools must locate activity records scattered across the filesystem and erase them irretrievably, while leaving the rest of the operating system intact. The technical challenge of finding and eliminating this data is far from trivial, given the complexity of modern computer operating systems, which are designed to preserve data rather than shed it. Yet published rigorous evaluations of these counter-forensic tools are limited.

This paper combines research findings from two rounds of testing, covering a total of 13 separate commercial counter-forensic tools. Successive versions of some tools were tested. The tests were designed to evaluate the tools' abilities to purge a range of activity records and other data representative of real-world computer use.

Although some tools performed markedly worse than others, all exhibited design or implementation shortfalls that lead to the disclosure of data with potential evidentiary significance. The counter-forensic software packages tested also left distinctive filesystem "fingerprints" from their usage.

These findings have led to the compilation of a reference resource for forensic analysts that covers the operational profiles of 19 tools and tool versions. Analysts may use this reference to associate a particular filesystem fingerprint with a counter-forensic too, and then use known operational weaknesses in that tool to guide their search for residual data. The process of screening disks for counter-forensic fingerprints can be automated to a large degree, and a new forensic utility, Aperio, has been developed to scan attached devices or images of filesystems for these traces. Aperio produces a detailed report, including the location of filesystem records it has flagged, facilitating the independent verification of its findings.

Weaknesses aside, it is important to note that most of the tools tested were capable of wiping beyond recovery (at least by conventional software-based forensic tools) the vast majority of the data they targeted. The resources outlined here can help speed and focus the analysis of suspected "disk-scrubbing", as well as demonstrate the use of counter-forensic tools where that activity has legal significance.

Background and related work

It is useful to differentiate the commercial counter-forensic software tested from counter-forensic utilities associated with computer system attackers. Computer intruders use tools such as log file cleaners and root-kits to hide their access to and activity on target systems. In contrast, the commercial tools evaluated in this report are aimed at *legitimate*

users of computer systems with various levels of technical proficiency. These tools are designed to eliminate specific activity records and user-designated files but leave the system otherwise complete and functional. All the tested tools were designed for the Microsoft Windows operating system, although at least one offers a variant for Apple platforms.

The focus and user profile of these commercial counter-forensic tools means they are less likely to be associated with cases of computer intrusions and pure computer crime, and more likely to be encountered in investigations seeking digital evidence of activity that is not solely computer related. Some specific examples are cited in the subsequent discussion on legal precedents.

Commercial counter-forensic tools' intended functionality may be broken down into two main areas:

- Locating relevant activity records on the system. This entails comprehensive, built-in knowledge of the data-handling behavior of the operating system and installed applications.
- Eradicating targeted data to thwart its recovery with standard forensic techniques. This typically entails overwriting the occupied data sectors on a disk with arbitrary values.

Failures in either functional area can lead to the disclosure of data that the tool's user sought to eliminate. Of the two areas, the second – data-wiping – has been more closely examined by researchers.

Methods have been developed to effectively destroy data on magnetic media, such as hard disk drives. One of the most frequently referenced standards in this area was produced by the U.S. Department of Defense in 1995 and recommends sanitizing magnetic media by overwriting repeatedly with specific patterns (DoD 5220.22-M). A year later, researcher Peter Gutmann published seminal research on recovering data from magnetic media using specialized tools and magnetic force microscopy. He also proposed a scheme for wiping data to thwart even a well-funded attacker, such as a government (Gutmann 1996).

Gutmann's threat scenario far exceeds the resources typically available at present to most forensic analysts. They rely on software tools to retrieve latent data from disks. Just overwriting the data once presents a major obstacle to recovery in these circumstances. As a result, forensic reviews of digital media often include an assessment of whether or not such counter-forensic tools were used, and it is has been suggested that these tools should be banned by corporate policy (Yasinsac and Manzano, 2001).

Related research includes a performance evaluation of counter-forensic tools by a team from University of Glamorgan that tested two tools, not identified by the researchers (Jones and Meyler, 2004). Research by forensic examiner Dan Jerger into artifacts from the installation and operation of counter-forensic software was presented at a closed-attendance conference last year (Jerger 2005).

On modern personal computer systems, two broad factors complicate the task of eliminating user files and activity records. One is the creation of arbitrary temporary files and cached data streams by common user applications, such as Microsoft Corp's Office suite or Internet Explorer web browser. Identifying and locating all the sensitive temporary data written to disk by user applications under varying circumstances is non-trivial. These temporary files are often deleted by the applications that created them, increasing the difficulty of locating the data subsequently in order to wipe it.

At the same time, operating systems use strategies to gain performance and data resilience that can propagate sensitive data onto arbitrary areas of storage media. These techniques include "swapping" data from RAM to a temporary "pagefile" on the disk to better manage system memory usage, and creating a file to store the contents of RAM

and system state information to support a "hibernation" function. Journaling file systems such as NTFS, ext3 and Reiser also record fractional changes to files in separate journal structures to allow filesystem records to be rebuilt more swiftly and consistently after a system crash (Shred manual pages, 2003).

The Analysis Results section of this paper discusses further challenges counter-forensic tools face in successfully locating and eliminating targeted data.

Legal precedent

Recent criminal and civil cases have tested the legal response to and treatment of these tools. One issue is the question of under what circumstances is demonstrated use of counter-forensic tools, by itself, an indication of consciousness of guilt or of intent to destroy evidence.

If the activity follows an order to preserve digital evidence, courts have ruled that the use of counter-forensic software is tantamount to the destruction of evidence and have sanctioned users (Kucala Enterprises v Auto Wax Co. 2003). In June 2005, Robert M. Johnson, a former publisher of Newsday and New York state education official, was charged with destruction of evidence for using counter-forensic software after learning he might be the target of a child pornography investigation (U.S. v Robert Johnson).

However, in the murder-kidnap case of Missouri v. Zacheriah Tripp, a state appeals court recently reinforced the view that simply using a wiping tool, in the absence of indications that material data was destroyed, is not suggestive of guilt (State of Missouri v Zacheriah S. Tripp, 2005).

In other cases, poorly used or improperly functioning data-wiping tools permitted the recovery of critical digital evidence (US v. H. Marc Watzman, 2003). And UK prosecutors have sought stiffer penalties for the use of a counter-forensic tool in recognition that evidence relevant to the severity of the crime was destroyed (O'Neill 2004).

Testing methodology

The test systems

Testing was conducted in two phases on two Intel Pentium desktop systems, one with 128MB of RAM and the other 256MB. Windows XP Professional Service Pack 1 was installed in the earlier phase testing on the first system. The same operating system with Service Pack 2 was installed on the second system. Both operating systems were installed on a 2.5GB partition formatted as an NTFS volume. Prior to the operating system's installation, the hard disk was overwritten with zeros to help ensure that previous artifacts on the media were not mistaken for data on the test system.

A principal user account was created with administrative privileges. This account was used for all subsequent activity on the systems.

In the Windows Internet Explorer (IE) browser, the privacy settings slider was moved to the "Low" setting for the Internet zone, one step below the default "Medium", to accelerate the collection of cookies. Form auto-completion was turned on. IE was set to delete browsing history records after five days, rather than the default 20. This was intentionally shorter than the intended test activity period to gauge the counter-forensic tools' abilities to eradicate history information that IE had already deleted. The size for IE's temporary cache of web pages, images and objects viewed was set to 25MB, about a third of its initially configured value.

Software installed for this study included: the Outlook, Word and Excel applications from Microsoft Office 2003; the Yahoo Instant Messenger client; the LimeWire P2P file-sharing client and the eDonkey P2P client; Adobe Acrobat Reader; and Macromedia Flash Player.

Activity record

Test activity on each system spanned about a week and sought to duplicate standard usage patterns for the installed applications. Wherever possible documents and records created during the test period were seeded with specific, repeated terms to help target subsequent searches for latent data.

Internet-related activity

Apart from browsing to a variety of web sites, this activity included:

- registering user accounts at websites such as the Washington Post, Hotmail, and Yahoo.
- joining P2P networks, enumerating available files and downloading several.
- saving HTML pages and linked components.
- conducting instant messaging chat sessions.
- retrieving and composing e-mail from Webmail accounts.
- transmitting and receiving POP3 e-mail via the Outlook Express and MS Outlook clients. E-mail was also moved between "folders" and deleted in these clients.
- downloading and installing software such as Acrobat Reader

Local activity

Available word processing clients, including Microsoft's WordPad, Notepad and Word, were used to create or edit several dozen documents in different formats. The process in Word was prolonged sufficiently to trigger the application's auto-save feature. This feature, which enables the recovery of "unsaved" work in the event of a power failure or application crash, saves a version of the document including changes to a temporary file that is deleted by Word if the document is subsequently closed normally. Images in various formats, principally JPG and GIF, were also saved, copied and edited. Files were also sporadically deleted or moved to the Recycle Bin. Discretionary file creation and manipulation occurred mainly in the test user's "My Documents" directory and its sub-directories

After several days of use, the eDonkey and Microsoft Excel applications were uninstalled from one of the test systems. One system was place in "hibernate" mode while an IM chat application was active. On the last day of the test period for each system files from every directory under the My Documents tree were moved to the Recycle Bin, and additional files from each directory were directly deleted.

Baseline filesystem image

At the end of the test activity period, the computers were shut down normally. Using the Helix bootable CD-ROM Linux distribution customized for forensic examinations, each computer was booted into a self-contained environment without mounting the hard drive's filesystems (Helix). A bit-for-bit image of the 2.5GB NTFS test partitions was made, using the Linux utility dd. After the imaging process, a checksum (using the MD5 hashing algorithm) of the imaged partition was compared to a checksum calculated on the original partition immediately prior to the image process. Matching checksums demonstrated that this provided a faithful copy of the full partition data, including deleted files and unallocated space. This image preserved the baseline configuration and activity record of the system before the installation of the counter-forensic tools to be tested.

Testing process

Configuration and use

The software packages evaluated were: Acronis Privacy Expert versions 7 & 8, Absolute Shield 3.42, CyberScrub Professional 3.5, Cyber Scrub Privacy Suite 4, Evidence Blaster 2005, Evidence Eliminator 5.058 build 9 & build 14, History Kill 2005, Privacy Eraser Pro 5.0, Privacy Guardian 4.0, R-Wipe & Clean 6.0, two versions of Secure Clean 4, TracksCleaner 3.0, Windows & Internet Cleaner Professional 3.60 [Privacy

Eraser Pro 5.0 appears to be a re-branded version of this software], two versions of Window Washer 5.5 and Window Washer 6 (Geiger 2005). Where the latest version was available under a fully functional trial license, this was used. Otherwise a license was purchased.

In each round of testing, the tools were installed into an identical operating environment created from the baseline filesystem image, allowing the performance of each tool to be tested against identical data and activity records. The counter-forensic software was configured and run, rebooting if recommended to complete the process. The system was then shut down normally from the Windows login prompt and booted into the Helix forensic environment described above. An MD5 hash was calculated for the Windows partition. A bit-for-bit image of the partition contents was created with dd, and the MD5 hash of the image file was compared to the pre-acquisition hash to verify the image was a faithful duplicate. A similarly validated copy of this image was used as a working copy for the analysis process.

Although configuration details varied from tool to tool, the set-up and use of the counter-forensic software followed a consistent approach:

- Each tool was configured to overwrite data targeted for deletion. A single overwriting pass was chosen, sufficient to obstruct recovery with standard software-based forensic applications. (See Figure 1 for an example of this step.)
- Most tools offered the option of obfuscating the names of files targeted for deletion, typically by renaming them with pseudo-random characters before deletion. This is designed to guard against disclosure of the names and types of files purged by the tools, since filesystem records about the deleted file can be retrieved even if the file's data contents are wiped. With this approach, a file named "Secret Ledger.xls" might be renamed to "XSFF443asajsa.csa" before deleting. This option was selected for each tool for which it was available.
- The tools were configured to eradicate Windows-maintained activity records such as browser history, document use history, the Internet Explorer file cache, recently used file lists, recent search terms, files in Windows temporary directories and stored cookies. Some of these records are contained in the Windows Registry database, some in other locations in the filesystem.
- Mail in selected Outlook Express and Outlook 2003 folders was targeted for secure deletion when the tool offered this option.
- Wiping the Windows pagefile, also referred to as the swap file, was selected in tools that offered this option. This contains data written from RAM to the hard disk, as the operating system seeks to juggle memory usage and performance.
- In tools that offered it, the wiping of unallocated, or free, space on the disk was selected. This entails overwriting disk data sectors that filesystem records do not show to be used by any active files.

- Each tool was used to purge the contents of the My Documents directory and subdirectories, and the contents of the Recycle Bin.

Not all the tools activated overwriting of data or unallocated space by default, although the tools' documentation often noted that such wiping is necessary to ensure that erased data

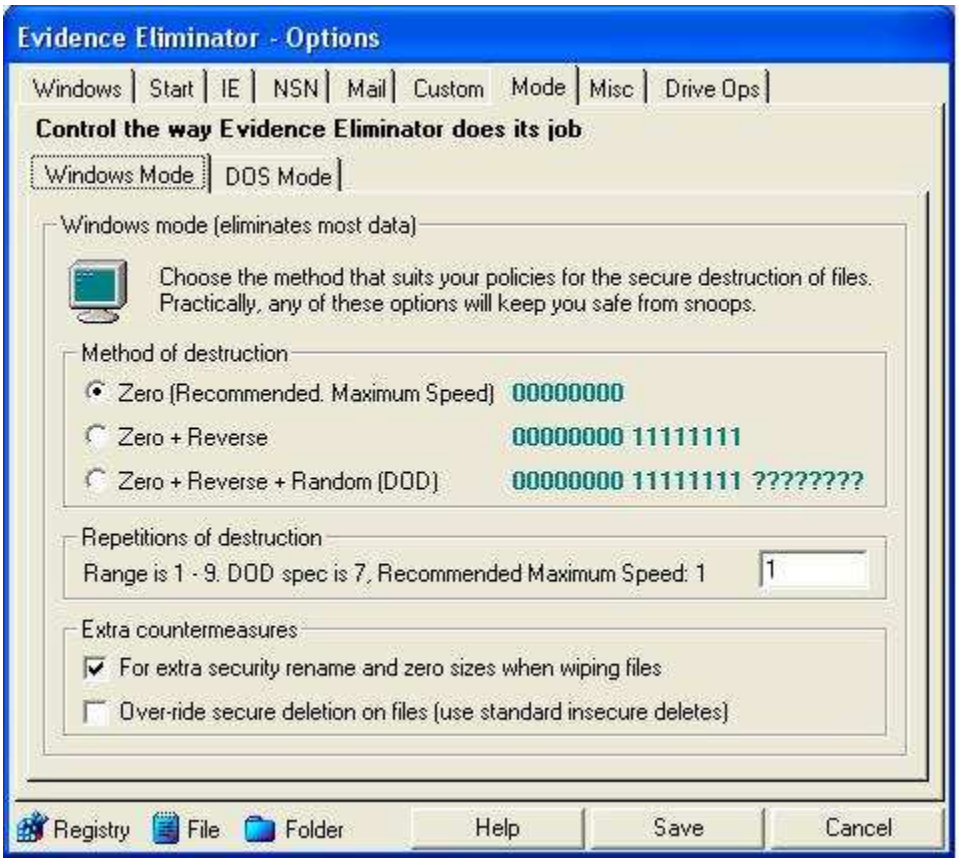

Figure 1: Configuring Evidence Eliminator

- Some tools offered the facility to eliminate activity records generated by third-party software, such as peer-to-peer clients, digital image studios and office productivity applications. Any that matched the installed applications described in the previous section were selected.
- The ability to wipe residual data in file slack space (the area between the end of data stored in a sector on the hard disk and the end of the sector) was also evaluated in some tools that offered this feature. (The earlier round of tests did not select this option).

are not recoverable. Under these default configurations, the forensic analyst's ability to recover data would greatly exceed what is generally reflected by this testing.

Analysis platform and tools

The main platform for analyzing the performance of the tools was the Forensic Tool Kit (FTK) version 1.50a-1.51 from AccessData. Like similar packages, FTK constructs its own map of disk space from the file system records, as distinct from the records that would be presented by the native operating system. Where filesystem metadata

still exist for deleted files (because they have not been overwritten or reallocated to new files), FTK can parse the information these records contain about the deleted files, including where on the disk those files' data was stored. FTK also processes unallocated, or "free," space on the disk for structured-data signatures and text content – and builds an index for later searching.

When file metadata has been obliterated, recovering data from the disk becomes more challenging, depending on the original data format. For most Microsoft Word documents, for example, much of the content exists in text format on the disk, and searching for a contained word or phrase can locate the deleted document's content on the disk. Many structured file formats, such as .jpg and .gif images or compressed Zip archives, can contain consistent sequences of code, or signatures. Using these location markers, the contents of the files can be reconstructed from contiguous unallocated disk space. This process is often termed "data carving."

Analysis results

All the counter-forensic tools failed to eradicate some potentially sensitive information – data specifically targeted for wiping by the user or records that contained information the tool was designed to eliminate. Some shortfalls were more serious than others. In four cases, the tools failed to wipe, or overwrite, most of the files slated for elimination, permitting the ready recovery of their original content.

In tools for which newer versions were evaluated in subsequent testing rounds, the later versions generally showed improvements, or at least no greater failures, in data elimination. In a couple of cases, some data leaks were plugged but new ones emerged. The most consistent improvement noted was in the purging of Registry data.

Although a similar testing framework was adopted in both series of tests, the testing environment was not identical and could possibly account for some variance from the first round of testing, subsequently referred to as CF-1 tests. In particular, Service Pack 2 was not installed on the CF-1 system, which also had half the RAM of the system used in the later round of tool trials.

Performance failures among the counter-forensic tool brands shared common features, reinforcing the view that some data-wiping tasks prove consistently difficult to achieve. One example is expunging 'resident' file data that is wholly contained in the NTFS Master File Table. Other filesystem structures, such as the journal file and pagefile, also proved challenging to purge. Performance-affecting software bugs occurred in many of the tools.

This said, these findings are only observed classes of failures and should be viewed as an indication of the performance characteristics of the tools tested, rather than as explicit prediction of data recoverability. Data recoverability will vary due to factors that range from hardware and software platform differences to more subtle influences on the computer system.

The results also suggest that the number of unique tools may be smaller than the number of vendors. Two evaluated wiping utilities were 're-branded' distributions of the same underlying engine -- Privacy Eraser from PrivacyEraser Computing Inc. and Windows & Internet Cleaner from NeoImagic Computing Inc. Apart from interface and configuration similarities, the operational characteristics of both matched closely.

Table 1 summarizes the areas of weakness and some representative examples of data recovery for a selected subset of the tools tested.

Table 1: Data Recovery Examples

	Wiping failures - 'free' space	Wiping failures - targeted files	Registry records missed	Activity files missed	Data recoverable from filesystem structures
Acronis Privacy Expert 8	Multi-paragraph text, fragment of browser History index, cached Web pages all recoverable.	File metadata not overwritten; Recycle Bin index file not wiped, allowing recovery of original metadata for bin contents; failed to delete designated mail from Outlook 2003.	Several keys under the ComDlg32 branch were intact, revealing recently used folders and documents; "Save As" key for Microsoft Office overlooked; IE download directory location.	Restore Point data, including user registry back-ups not eliminated; omitted a Windows prefetch folder component that contained path info and names of wiped files; recent Office document shortcuts missed; IE browsing cache index file only partially wiped.	Cookies, small images resident in MFT recoverable; file name data disclosed in NTFS journal.
Absolute Shield 3.42	Multi-paragraph text, cached Webmail page views recoverable.	User cannot designate arbitrary files for wiping; metadata not overwritten.	Yahoo profile chat partner disclosure; "Save As" key for Microsoft Office overlooked; ShellNoRoam/ BagMRU key revealed file names.	Restore Point data, including user registry back-ups not eliminated; Windows prefetch folder intact, disclosing path and names, other data, for wiped files.	Cookies, small images resident in MFT recoverable; file name data disclosed in NTFS journal; directory index files for some folders contained entries disclosing wiped file names; pagefile contained multi-paragraph text from deleted documents.

	Wiping failures - 'free' space	Wiping failures - targeted files	Registry records missed	Activity files missed	Data recoverable from filesystem structures
Cyber Scrub Privacy Suite 4	Multi-paragraph text, fragments of Web pages recoverable; file slack incompletely wiped, revealing fragments of documents, web data.	Yahoo IM client log renamed but not deleted.	RecentDocs key in user hive listed recently used documents, broken down by file type.	IE History, IE cache and IE cookie index files all were missed; Windows Media Player resource file with content URL missed; Windows prefetch folder intact, disclosing path and names, other data, for wiped files.	Wiped file name data disclosed in NTFS journal and in a few unallocated MFT entries.
Evidence Blaster 2005	Did not feature wiping of unallocated space; extensive data recovery possible.	Deleted but failed to wipe user-specified documents and files under My Docs folder; also failed to wipe contents of Recycle Bin, IE cache files and cookies; metadata not obfuscated.	ComDlg32 branch was intact, revealing recently used folders and documents; recent search terms also not wiped; WordPad recent file key intact; "Save As" key for Microsoft Office overlooked.	IE History index file incompletely purged, leaving some record elements; IE browsing cache index file was also only partially cleaned; Yahoo IM client log missed; Windows Media Player resource file with content URL missed.	Pagefile contained multi-paragraph text from deleted documents and web content; directory index files for some folders contained entries disclosing wiped file names.
Evidence Eliminator 5.058 b14	A small number of references to wiped file names and paths were recoverable.	__eetemp directory in filesystem root contained undeleted copies of prefetch folder files, as well as IE History, cache and cookie index files.	"Save As" keys for Microsoft Office applications missed; ShellNoRoam/ BagMRU key revealed file names.	Windows Media Player resource file with content URL missed.	Directory index files for some folders contained entries disclosing wiped file names.

	Wiping failures - 'free' space	Wiping failures - targeted files	Registry records missed	Activity files missed	Data recoverable from filesystem structures
History Kill 2005	Did not feature wiping of unallocated space; extensive data recovery possible.	User cannot designate arbitrary files for wiping; deleted but failed to wipe many of the files targeted for elimination, including contents of Recycle Bin and IE cache; browser History files deleted but not wiped.	ComDlg32 branch was largely untouched, revealing recently used folders and documents; Acrobat recent file list remained; WordPad recent file key intact; "Save As" key for Microsoft Office overlooked.	Restore Point data, including user registry back-ups not eliminated; Windows prefetch folder intact, disclosing path and names, other data, for wiped files; shortcut files to recent Office documents overlooked.	Cookies, small images resident in MFT recoverable; wiped file name data disclosed in NTFS journal; pagefile contained multi-paragraph text from deleted documents, web pages.
Privacy Eraser Pro 5.0	Multi-paragraph text and fragments of Web pages recoverable; file slack incompletely wiped, revealing fragments of documents, web data.	Restore Point files deleted but not wiped, allowing recovery; same true of Windows prefetch folder contents and Recycle Bin index file.	"Save As" key for Microsoft Office overlooked; ShellNoRoam/ BagMRU key revealed file names.	IE browsing cache index file was also only partially cleaned; Windows Media Player resource file with content URL was missed; Yahoo IM client log missed.	Wiped file name data disclosed in NTFS journal and in a few unallocated MFT entries.

	Wiping failures - 'free' space	Wiping failures - targeted files	Registry records missed	Activity files missed	Data recoverable from filesystem structures
Privacy Guardian 4.0	Multi-paragraph text and fragments of Web pages recoverable.	Wiping failed for most of the user-specified files under the My Docs folder tree – some were also left undeleted; IE cache index only partially purged.	"Save As" key for Microsoft Office overlooked.	Windows prefetch folder intact, disclosing path and names, other data, for wiped files; missed MS Office shortcuts to recently used files and Yahoo IM client log; IE cookie index file overlooked.	Wiped file name data disclosed in NTFS journal and in a few unallocated MFT entries; pagefile contained multi-paragraph text from deleted documents, fragments of web pages and IM chat sessions; cookies and images small enough to be resident in the MFT were recoverable.
Secure Clean 4	A small number of references to wiped file names and paths were recoverable; slack space for one file contained fragments of cached web page.	Failed to purge deleted mail stores in both Outlook Express and Outlook 2003.	"Save As" key for Microsoft Office overlooked; ShellNoRoam/ BagMRU key revealed file names.	Windows prefetch folder intact, disclosing path and names, other data, for wiped files;	Wiped file name data disclosed in NTFS journal and in a few unallocated MFT entries.

	Wiping failures - 'free' space	Wiping failures - targeted files	Registry records missed	Activity files missed	Data recoverable from filesystem structures
TracksCleaner 3.0	Function to wipe unallocated space froze without progress repeatedly on test system; slack space contained numerous fragments of targeted documents and web content.	Failed to wipe most user-targeted files, and Recycle Bin contents, conventionally deleting instead; also failed to wipe program-selected files, including IE cache contents, files in the system Temp folder and browser History files.	ComDlg32 tree largely untouched, revealing recently used documents; Acrobat recent file list remained; Explorer Recent Docs key intact; "Save As" key for Microsoft Office overlooked.	Restore Point data, including user registry back-ups not eliminated; Windows prefetch folder intact, disclosing path and names, other data, for wiped files; IE cache index file left untouched.	Wiped file name data disclosed in NTFS journal and in a few unallocated MFT entries; pagefile contained numerous web page fragments and multi-paragraph text from deleted documents, and IM chat sessions;
Window Washer 6	Multi-paragraph text, fragments of Web pages recoverable; file slack incompletely wiped, revealing fragments of documents, web data.	Deleted message fragment recovered from Outlook Express mail store; browser History file contained a few references to activity prior to wiping tool's operation.	"Save As" key for Microsoft Office overlooked; ShellNoRoam/ BagMRU key revealed file names.	Restore Point data, including user registry back-ups not eliminated; Windows prefetch folder intact, disclosing path and names, other data, for wiped files; IE cache index file left untouched; missed MS Office shortcuts to recently used files and Yahoo IM client log.	Pagefile contained numerous web page fragments and multi-paragraph tracts from deleted documents, and IM chat sessions; wiped file name data disclosed in NTFS journal and in a few unallocated MFT entries; directory index files for some folders contained entries disclosing wiped file names.

Analysis discussion

This section provides a narrative description of the classes of failures observed.

Incomplete wiping of unallocated space

Searching unallocated disk space – areas of the disk registered as unused in the filesystem index – for occurrences of the terms seeded during testing and for structured file signatures recovered substantial data from all but five of the tools tested. Files and fragments of files conventionally deleted by the user would be expected to exist in unallocated space, along with content from the plethora of temporary files routinely created and deleted by the operating system and applications. Some tools did not provide the option to wipe unallocated space. These cases, and where wiping failures were extensive, provide significant scope for the recovery of latent data, including complete files.

For example, Microsoft Word creates temporary copies of documents to record uncommitted changes to aid in recovering from a crash. The copy is automatically deleted when the Word document is closed normally – but because the deletion operation only affects the file's index record, what this really means is there is no longer a convenient way to locate the document contents on the disk in order to overwrite it. Forensic tools designed to find exactly such orphaned information on the disk can still rebuild the document. Other deleted copies of the data may have been scattered elsewhere on the disk, created as temporary copies during the download process or by virus-scanning software. The scope of data recoverable from unallocated space has been well demonstrated by researchers examining disks bought second-hand (Garfinkel and Shelat, 2003).

Failure to erase targeted user, system files

All the counter-forensic tools missed some records created by the operating system or user applications that contained sensitive information. Some of this data disclosure resulted from failures to fully overwrite files targeted for deletion; in other cases, there was no evidence the tool attempted to delete the relevant data.

In general, the greater the range of third-party applications any tool attempted to incorporate as data-wiping targets, the more likely it was to encounter failures in that coverage. Some tools attempt to track more than a hundred third-party applications such as photo-editing suites and media players. Research suggests the resources required to track the location of activity records generated by so great a range of evolving software and OS interactions heavily taxes the resources available to counter-forensic tool developers (Geiger and Cranor, 2005).

Similarly, changes in the base operating system's functionality created data leaks for most of the tested tools. For example, most tools ignored the prefetch folder introduced in Windows XP (Windows XP Development Kit 2004). The folder's contents, which are used to speed the loading of files frequently accessed by the system or user, include files that detail the full path and names of many of the files in wiped directories.

Because of another function introduced with Windows XP, the tools' ability to purge usage data stored in the Windows Registry proved moot in many cases. Most tools overlooked back-up copies of the user registry stored as part of Windows XP's creation of "restore points" for the system. These restore points, triggered on schedule or by configuration changes, record system configuration information, often including copies of user registry files. The back-up registry copies contained *essentially all* the records the tools sought to delete from the user registry. In fact, the installation of the wiping tools frequently triggered a restore point back-up of key configuration files, including a copy of the user's registry hive just before the use of the tool.

Registry usage records missed

As noted, the newer versions of previously evaluated counter-forensic tools generally improved their ability to clean activity records in the Registry, a centralized database structure used by the operating system and applications to hold configuration information, license data and a wide array of other details about the system and installed software.

Still, some Registry entries introduced by Microsoft Office 2003 escaped purging by a few of the tools and most missed keys under the ShellNoRoam/BagMRU branch of the user registry hive, which contained references to a few wiped files' names and locations.

Data recoverable from special filesystem structures

All the tested tools encountered problems eradicating some data from special filesystem structures. The operating system usually curtails access to these structures by user applications because they are critical to the filesystem's integrity.

Tiny text files, such as browser cookies, and some small .gif images cached from web activity were recoverable from the NTFS Master File Table (MFT). The MFT, the main index to information about files on the filesystem, can also contain a file's data if it amounts to roughly 700 bytes or less (Carrier 2005, p. 283). This "resident" data exists as a component within the MFT special file structure, and wiping this space proved problematic for the tools.

Small files and fragments of larger files were similarly recoverable from the NTFS filesystem journal after most tools were run. The journal file stores partial changes to files before they are written to the filesystem to make recovering from a crash simpler and faster. A number of tools also failed to eliminate sensitive data from the pagefile. As another special system file, this might have presented wiping problems for the counter-forensic tools, although Windows XP offers a built-in facility to overwrite the pagefile on system shutdown.

Fingerprinting counter-forensic tools

Most of the tested tools left distinctive signatures of their activity that could be used to postulate the tool's use even if no evidence of the software's installation was recovered. (This might occur, for example, if a tool installed on a separate partition or physical disk is used to delete data on another.) The patterns they created in the filesystem records would not be expected to occur during typical computer activity.

The most common distinguishing pattern created by the use of the tested tools' was their technique for mangling metadata about files they wiped. In particular, all the tested tools that renamed the files they sought to wipe adopted differing strategies for generating new file names. Most of the counter-forensic packages tested offered to rename wiped files (and often alter other data, such as file size and creation date) in order to minimize the information that can be gleaned by examining the metadata for deleted files. The operating system controls the allocation of MFT entries that store this metadata on NTFS filesystems. The entries are not cleared when a file is erased or wiped, only when they are reallocated to a new file entry. Figure 2 outlines the basic data fields contained in a file entry in the MFT.

Other salient patterns left in MFT records by the tested tools included: whether the file size was set to zero and pointers to the file data sectors were cleared;

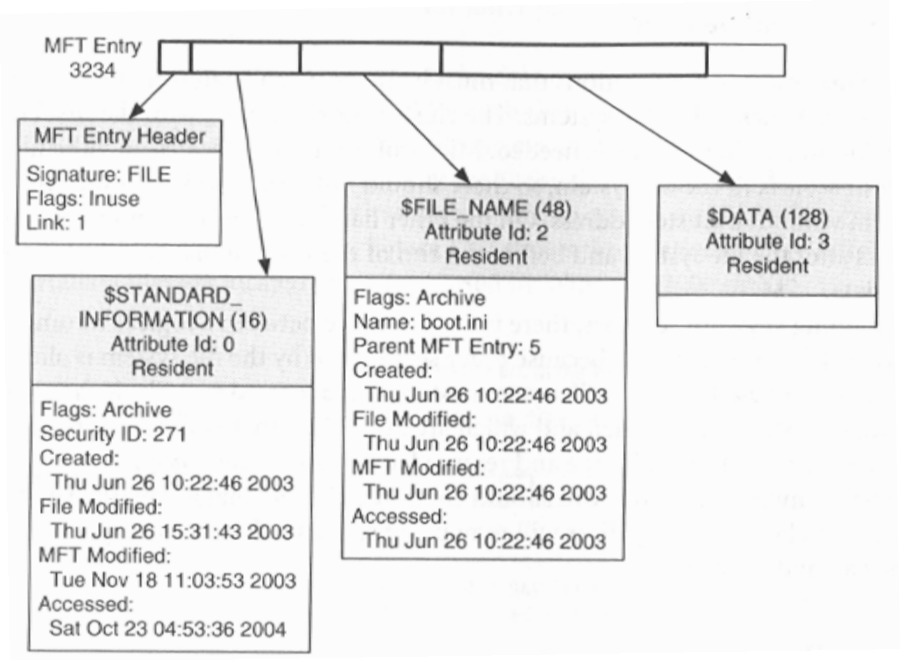

Figure 2: Structure of an MFT record (Carrier 2005)

whether file date records were altered and how; and whether Alternate Data Streams were subject to renaming. With some tools, operational signatures were also identifiable on other filesystem structures. These included specific patterns (other than simply zeroes) used in wiping disk data sectors and the creation of consistently named files and/or directories during the wiping process.

In the case of the minority of tools that generated less definitive signatures from their file renaming approach, combining other MFT and filesystem traces was usually sufficient to specify a distinguishing fingerprint. Table 2 provides a range of signature characteristics for a selected subset of the counter-forensic tools tested.

In several cases, the file renaming strategy adopted uses an incrementing counter approach to generate unique file names. So, recovery of an MFT entry with this pattern can provide a lower bound to the number of files wiped during that session. Often specific file time fields were observed to be set to the date and time of the session as well, providing another dimension of potentially valuable data about the activity. Given the earlier noted legal precedents about the use of these tools, the presence of such signatures might have probative value in some cases.

Limitations

A number of important limitations and considerations apply to the use of these fingerprints to postulate the operation of a counter-forensic tool and make assertions about terms of its operation. A basic requirement is that other potential sources of these signatures, including their intentional fabrication, be considered. Although it seems highly unlikely that the routine use of other software would generate the full constellation of patterns in any well-defined signature, some component patterns may be created by other activity. Context, such as the number of these records found, is also relevant.

For example, while programs other than Privacy Eraser Pro may use the Windows GUID-generating API to produce unique file names, fewer would use the complete value and a consistent file extension of ".tmp". And it would be dramatically less likely for any to create a significant number of zero-length files with that full naming scheme.

The following additional limitations are noted:

- The signatures specified were solely for NTFS volumes. Although NTFS is replacing the FAT filesystem format on Windows operating systems, FAT-formatted filesystems are still common. The FAT filesystem differs markedly from NTFS, and tool signatures are likely to vary also.
- Most of the tools that offered metadata scrambling allowed it to be disabled. Although this would allow the recovery of original metadata for wiped files, key elements of the tool signatures would be absent.
- Other configuration changes, such as specifying the character used for overwriting data sectors (permissible in a few tools), could alter other constituents of the fingerprints.
- Although signatures have so far uniquely identified the tools for which they have been specified (or the underlying engine in the case of Privacy Eraser Pro and Windows & Internet Cleaner), tools that have not been tested may share the signature of those that have.
- A version change in one tool was accompanied by a sharp divergence in its fingerprint. Consequently, unknown wiping fingerprints need not imply the use of a tool other than those tested, but perhaps only a different version.

Installation artifacts

Any opinion about the use of these counter-forensic tools may be bolstered by the identification of files and other artifacts associated with the installation of the programs themselves. Practitioners and forensic utility vendors have compiled lists of Registry keys and files (and their hash values) associated with the installation of many of the counter-forensic tools tested (Jerger 2005; Brown 2005).

The two approaches to identifying these tools are complementary. Although installation artifacts can demonstrate the presence – current or past – of this software, the operational signature generated by this category of tools can better demonstrate their actual use.

Table 2: Selected Tool Signatures

Utility	Operational signature
Acronis Privacy Expert 8	Privacy Expert left filesystem metadata intact but overwrote associated data sectors with zeroes. By itself, this does not comprise a conclusive filesystem-level signature because a similar pattern may be present after the use of other tools – if metadata scrambling features are not selected on those tools. However, this approach leaves metadata recoverable that may provide useful information to investigator.
Absolute Shield 3.42	Absolute Shield created files of the form SSF[somevalue].tmp when it renamed and overwrote targeted files. Files point to space that has been wiped either with x00 or xFF characters. [somevalue] is one or more characters, each in the hexadecimal value range (0-9 and A-F). The value appears to be linearly incremented as files are wiped. Example: SSF9A1.tmp
Cyber Scrub Privacy Suite 4	Wiped files were renamed with pseudo-random combinations of capital letters of varying lengths and varying three-letter extensions. Example: WEFOPSDFSQ.JKV. File data length is recorded as zero. A deleted, temporary file with the extension ".wip" is created in the volume's root directory. In addition, a more complex fingerprint appears to result from steps apparently taken to overwrite MFT entries. Filesystem records showed a large number of deleted files with the name x.tmp, where x is a single-digit number. These were nested in directories also named with a single-digit number under a directory in the volume root called (in this instance) Erase5A4.tmp/. Example: C:/Erase5A4.tmp/3/2/1/5/4.tmp. All of these deleted files point to random-looking data small enough to be MFT resident. Another recurring feature is the existence of a large number of Alternate Data Stream entries named just "a" for deleted and wiped files.
Evidence Blaster 2005	The tool failed to wipe targeted data and did not rename files or scramble metadata, so its operational signature at the filesystem level is harder to define. However, the level of data recovery possible may make finding signature data less important in many cases. An alternate candidate for determining Evidence Blaster's use is its treatment of the IE History index file. Evidence Blaster only partially cleans the IE History index file containing user browsing records. While URL and user names were removed from records, other fields such as access times, visit count and page title were left.
Evidence Eliminator 5.058 b14	Wiped files were renamed with 243 characters with no filename extensions. All except the first 10 characters are pseudo-random combinations of lowercase letters. The first 10 characters are numbers that appear to increment by one for every file wiped. Example: 0000002825wtkdvjiiugvwgveodruvlmdptxgpgfyrqnxpxyjajk qrienrnebnzhoshuyfzhdvzvvvveszlikswlhqpwbetowmznlvzquveyvhkrk cidsmpgpjrxjgpzaxcffvdxynlxiikdnhgachijkuajmdfdcvxbupesrwdyykqf ckndbqwittwnyfmtcesftoxtyrnfdwwoblkpcvzwseokhydmcvtvodbrwyv vmewuoge The creation of the __eetemp directory in the filesystem root (see notes) may also be considered signature behavior.

Utility	Operational signature
History Kill 2005	The tool failed to wipe targeted data and did not rename files or scramble metadata, so its operational signature at the filesystem level is harder to define. However, the level of data recovery possible may make finding signature data less important in many cases.
Privacy Eraser Pro 5.0	In overwriting targeted file data, Privacy Eraser Pro renamed files and truncates their length to zero prior to deleting them. The renamed files were named in the form XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX.tmp, where the X's are replaced with capital letters and numbers in the hexadecimal value range. All use .tmp as the file extension. Example: 4B282BCB-C34D-4147-ACFA-645F3D524B8D.tmp This signature is identical to that of Windows & Internet Cleaner 3.60, which appears to share the same engine and features. The value chosen for the renaming scheme appears generated by Windows' GUID-creation function. In operating systems prior to Windows 2000, the computer interface's MAC address was incorporated as part of the GUID, which could more concretely tie the activity to a particular system (Gutmann 2005).
Privacy Guardian 4.0	Privacy Guardian did not munge the metadata entries for files it deletes, and so left a less defined operational signature. However, the way the program wiped system Restore Point files might provide evidence of the tools use -- all were overwritten with zeroes before being deleted, in contrast to its handling of other files. The data areas for other successfully wiped files appear to have been overwritten with random values, which may be flagged as high-entropy areas by analysis tools.
Secure Clean 4 on XP-SP2	SecureClean renamed files during its name and metadata scrambling operations. The file names take the form of "SCxxxxxx.T~P", where the 'x's stand for a six-digit number that seems to increment linearly for every file wiped. Example: SC000043.T~P These files' MAC times were set at the time of wiping, and the file size set to zero. In addition, to the filesystem signature noted above, this test identified other operational fingerprints. In Secure Clean's overwriting of Restore Point files, the names and other metadata of these files were unchanged -- although the corresponding data sectors were overwritten with zeroes. The MAC times for files in the Restore Point directories did not appear to be altered. SecureClean also left metadata pointing to a deleted folder in the root directory named "sctemp", which contained deleted files with names in the form "AF~Sxxxx.T~P", where xxxx was a four-digit number. Judging by file size and other remaining metadata records, these files may have been created during the overwriting of unallocated space.
TracksCleaner 3.0	The tool did not rename files or scramble metadata, and its overwriting was inconsistent, so its operational signature at the filesystem level is harder to define. However, the level of data recovery possible as a result may make finding signature data less important in many cases.
Window Washer 6	Targeted files renamed with varying lowercase letters for both the filename and a three-letter extension. The length of filename also varied. Example: fpubhmrwbgkpuydin.ydh. Values used to overwrite data sectors varied from file to file, but this character is repeated for the full space allocated to the file.

Counter-forensic analysis resources

To facilitate the application of this research in the analysis of suspected counter-forensic activity, two resources have been created for forensic examiners.

Web-based reference

One resource is a Web-based reference guide to performance reports on tested tools and their signature specifications (Figure 3). This resource also includes suggestions to maximize data recovery depending on the particular tool's characteristics. Notes and observations on the tools are accompanied by a detailed forensic report on its operational test, along with recovered artifacts.

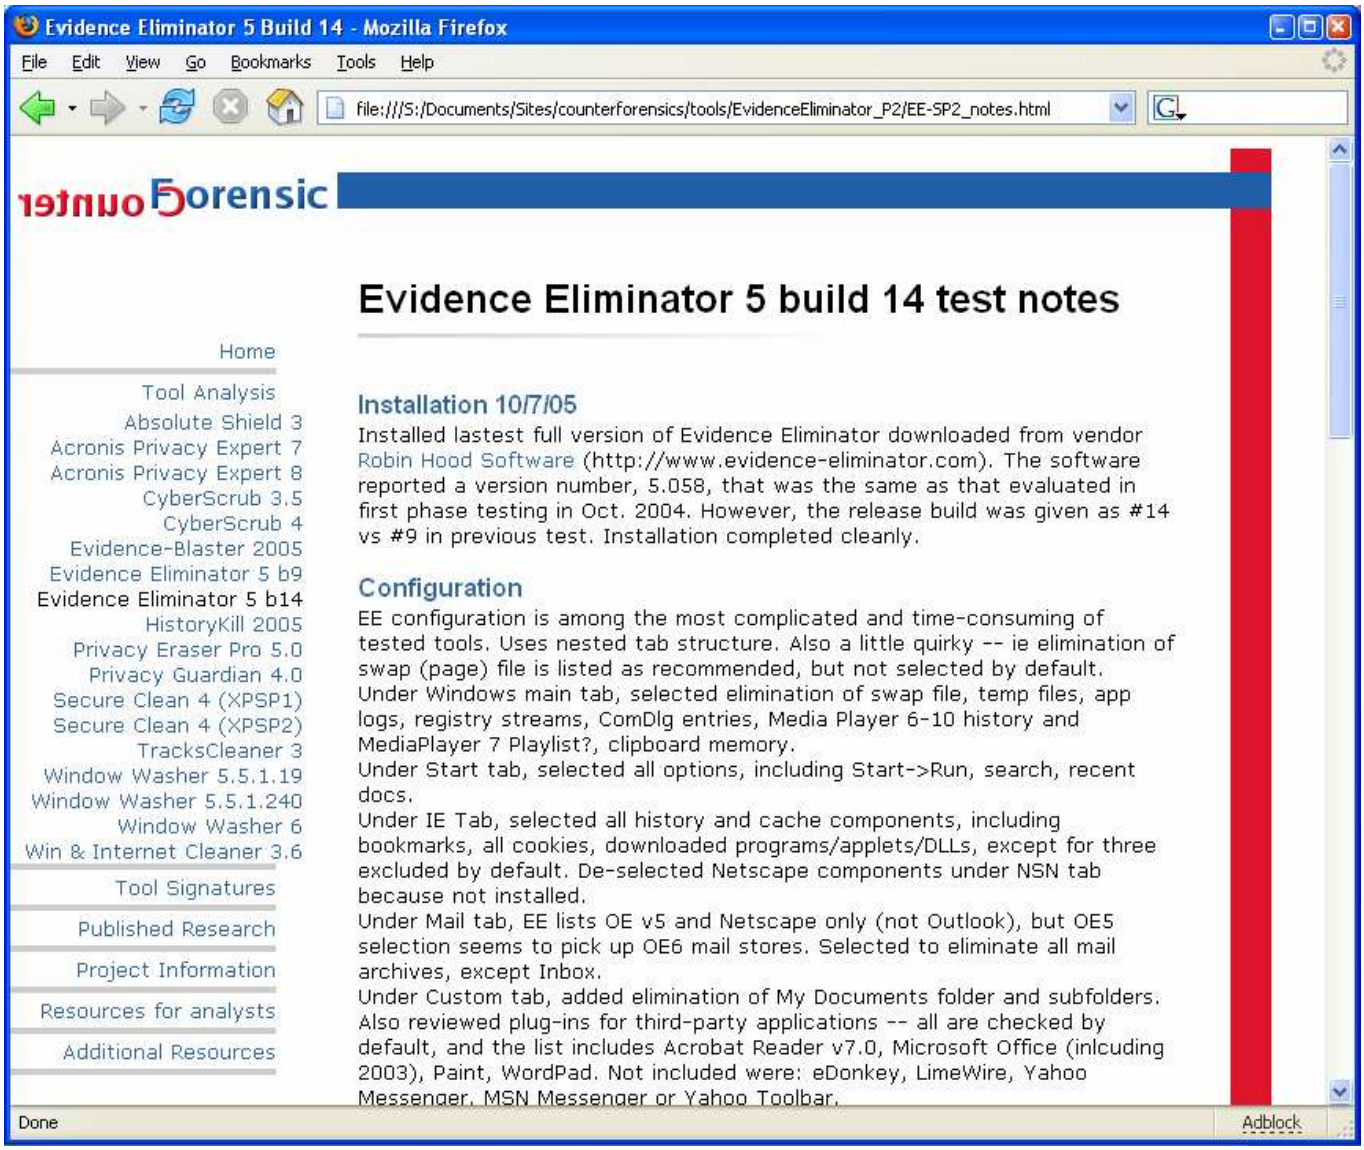

Figure 3: Web reference

Aperio

A second resource is Aperio, a forensic utility that searches for signatures of tested counter-forensic tools. This program, written in C, uses the libraries of the open-source Linux-NTFS filesystem project to help locate and parse NTFS filesystem structures, such as the MFT. It is designed for compatibility with CD-bootable Linux distributions often used in forensic analysis, and should compile on most Linux platforms.

Aperio creates a data structure that stores signature parameters and data about the tools to which they apply. Apart from tool name and version information, the structure fields map to the timestamp, file size, name and non-resident data patterns for an MFT entry. Signature specifications for each of these fields are read from a configuration file. The configuration file is readily customizable to add new signatures as they are specified. Extension of the program's signature logic can be achieved through changes to the data structure and/or logic functions.

The utility uses C's Regular Expression library to match patterns in the name fields of deleted MFT records and from associated data sectors. Aperio will process filesystems presented to it as imaged files or as devices. Its actions are read-only and will not alter the filesystem it scans. Aperio's verbose output specifies the location of records it flags and reproduces their MFT entry data, simplifying validation and debugging of its findings.

A functional version of Aperio that fully implements regular expression matching within the MFT name field and limited logic on other fields has been produced. Both Aperio and the reference resource described above are being utilized on a trial basis by law enforcement and government agencies. Wider distribution is under consideration.

Acknowledgements

The U.S. Department of Homeland Security has supported this research and its future application. My colleague, Richard Nolan of CERT, provided resources and support necessary to complete this work in a relatively compressed period. He also offered invaluable insight and perspective from a highly relevant professional background.
My initial research in this area was encouraged by Lorrie Faith Cranor, at the Institute of Software Research, International of Carnegie Mellon University. We also collaborated on a paper that focused on the privacy implications of these findings.
For advice, suggestions and useful criticism, sincere thanks are also due to Simson L. Garfinkel and Peter Gutmann.

Future work

The value of the described reference resources and Aperio could be increased by extending the collection of tool performance and signature data. This would include untested tools, new (and older) versions of tested tools and an examination of the impact of system configuration differences on tool behavior. Performing parallel testing on FAT filesystems would be similarly worthwhile.

Work is underway to extend the logic and discriminatory ability of Aperio with the aim of maximizing reliability and coverage.

References

Brown, Christopher L.T. Conversation with the author on the compilation of hash sets of counter-forensic tool installation artifacts by Technology Pathways LLC, of which he is CTO. August 2005.

Carrier, Brian. *File System Forensic Analysis.* Addison Wesley, 2005.

Geiger, Matthew. "Evaluating Commercial Counter-Forensic Tools", in proceedings of the *5th Annual Digital Forensic Research Workshop*, New Orleans, Louisiana, Aug 17-19, 2005. http://www.dfrws.org/2005/proceedings/geiger_couterforensics.pdf

Geiger, Matthew and Lorrie Cranor. "Counter-Forensic Privacy Tools: A Forensic Evaluation." Technical Report, Institute for Software Research Intl, Carnegie Mellon University, June 27, 2005. http://reports-archive.adm.cs.cmu.edu/anon/isri2005/CMU-ISRI-05-119.pdf

Garfinkel, Simson L. and Abhi Shelat. "Remembrance of Data Passed: A Study of Disk Sanitization Practices." IEEE Security & Privacy magazine, January/February 2003.

Gutmann, Peter. "Secure Deletion of Data from Magnetic and Solid-State Memory." First published in the *Sixth USENIX Security Symposium Proceedings*, San Jose, California, July 22-25, 1996. http://www.cs.auckland.ac.nz/~pgut001/pubs/secure_del.html

Gutmann, Peter. Private e-mail correspondence. Nov. 28-29, 2005.

Helix Incident Response & Computer Forensics Live CD. Available from http://efense.com/helix/index.php.

Hopper, Ian D. "Enron's Electronic Clues: Computer Scientists Seek to Recover 'Deleted' Files." Associated Press, Jan. 16, 2002. Viewed at: http://abcnews.go.com/sections/scitech/DailyNews/enronPCfiles020116_wire.html

Jerger, Daniel S. "Artifacts of Popular Data Deletion Utilities", presentation at HTCIA 2005 International Conference. Monterey, California, Aug 29-31, 2005. Not publicly available.

Jones, Andy and Christopher Meyler. "What evidence is left after disk cleaners?" Digital Investigation: The International Journal of Digital Forensics & Incident Response, ISSN: 1742-2876, Vol 1, No 3 Sep 2004: 183-188.

Kucala Enterprises v Auto Wax Co. (2003). Judgment in case # 02C1403, United States District Court, Northern District of Illinois. Available as Case No. 1403 - Doc. No. 127 from http://www.ilnd.uscourts.gov/racer2/.

Windows XP Development Kit. "Memory Management Enhancements." Online article on Microsoft Developer Network, Sep 2, 2004. http://msdn.microsoft.com/library/default.asp?url=/library/en-us/appendix/hh/appendix/enhancements5_0eecebea-e58b-4c95-8520-9b1dc2bc6196.xml.asp

O'Neill, Sean. "Court battle on software that destroys cases against paedophiles." The Times of London, Dec. 3, 2004. http://www.timesonline.co.uk/

Shred manual pages. A component of the Linux coreutils package v 4.5.3, November 2003. Documentation available as part of the coreutils distribution and at http://techpubs.sgi.com/library/tpl/cgi-bin/getdoc.cgi?coll=linux&db=man&fname=/usr/share/catman/man1/shred.1.html&srch=shred

State of Missouri v Zacheriah S. Tripp. Appeal ruling on case # WD63005 in the Missouri Court of Appeals Western District, June 7, 2005. http://www.courts.mo.gov/courts/pubopinions.nsf/0/a146c42d2b4a0e6486257018005e0edd?OpenDocument

United States v. H. Marc Watzman (2003). Indictment in United States District Court, Northern District of Illinois, Eastern Division. http://www.usdoj.gov/usao/iln/indict/2003/watzman.pdf
See also http://www.kansas.com/mld/kansas/news/7119391.htm for a report of the case.

U.S. Department of Defense "Standard 5220.22-M: National Industrial Security Program Operating Manual" (January 1995), Chapter 8. http://www.dss.mil/isec/chapter8.htm

U.S. v Robert Johnson. Three-count indictment filed in U.S. District Court, Southern District of New York, June 28, 2005.

Yasinsac, Alec and Yanet Manzano. "Policies to Enhance Computer and Network Forensics." *Proceedings of the 2001 IEEE Workshop on Information Assurance and Security*, United States Military Academy, West Point, New York, 5-6 June, 2001. http://www.itoc.usma.edu/Workshop/2001/Authors/Submitted_Abstracts/paperW2B3(37).pdf