

DFRWS USA 2016 — Proceedings of the 16th Annual USA Digital Forensics Research Conference

Time is on my side: Steganography in filesystem metadata



Sebastian Neuner^{a,*}, Artemios G. Voyiatzis^a, Martin Schmiedecker^a,
Stefan Brunthaler^a, Stefan Katzenbeisser^b, Edgar R. Weippl^a

^a SBA Research, Vienna, Austria

^b Technische Universität Darmstadt, Germany

A B S T R A C T

Keywords:

Digital forensics
Data hiding
Steganography
Storage forensics
File system forensics
Real-world data corpus

We propose and explore the applicability of file timestamps as a steganographic channel. We identify an information gap between storage and usage of timestamps in modern operating systems that use high-precision timers. Building on this, we describe a layered design of a steganographic system that offers stealthiness, robustness, and wide applicability. The proposed design is evaluated through theoretical, evidence-based, and experimental analysis for the case of NTFS using datasets comprising millions of files. We report a proof-of-concept implementation and confirm that the embedded information is indistinguishable from that of a normal filesystem use. Finally, we discuss the digital forensics analysis implications of this new information-hiding technique.

© 2016 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Introduction

The need for protected information exchange and storage in the digital world is constantly increasing. Cryptographic techniques can provide information confidentiality, authenticity, and integrity. However, they do leave evidence of the information exchange.

Steganographic techniques are able to hide the existence of information passing through communication channels or resting in storage media for later access. These techniques are useful in a wide range of real-world scenarios, including but not limited to: circumventing censorship and restrictions imposed by governments and other adversaries (Akgül and Kırıldoğ; Anderson, 2012), assisting whistleblowers when disclosing documents (Greenwald, 2014), and supporting businesses to protect

strategic corporate information during transmission (Cox et al., 2007).

Numerous steganographic techniques have been proposed and analyzed in the research literature (Zielińska et al., 2014). The analysis focuses on criteria such as the achieved secrecy on specific application scenarios, the steganographic channel capacity, and the information channel utilization.

Storage or format-oriented steganographic techniques hide information in logical channels by utilizing redundant or unused fields in format specifications. This includes, among others, the master boot record (MBR) of non-bootable hard disks and the unused disk space caused by the misalignment of hard disk sector size and file size (Khan et al., 2011).

Modern filesystems support a wealth of operations that span beyond the primitive of mapping files into sequences of hard disk sectors. The filesystem specifications define additional data structures (i.e., “metadata”) to describe information like the owner, the access permissions, and the date and time when important file events took place.

In this paper, we propose and explore, to the best of our knowledge for the first time in literature, the applicability

* Corresponding author.

E-mail addresses: sneuner@sba-research.org (S. Neuner), avoyiatzis@sba-research.org (A.G. Voyiatzis), mschmiedecker@sba-research.org (M. Schmiedecker), sbrunthaler@sba-research.org (S. Brunthaler), katzenbeisser@seceng.informatik.tu-darmstadt.de (S. Katzenbeisser), eweippl@sba-research.org (E.R. Weippl).

of *filesystem timestamps* as a steganographic channel. More specifically, we make the following contributions:

1. We analyze the granularity of the timestamps that modern filesystems implement, and we evaluate their applicability for steganographic applications.
2. We propose the use of timestamps as a means to hide information in NTFS and other filesystems with sub-second timestamp granularity.
3. We describe a system design and a proof-of-concept implementation that support different levels of possible capacity to securely hide data on NTFS volumes.
4. We validate the proposed system using real-world and synthetic datasets, and we show that the embedded steganographic information cannot be distinguished from the information produced by normal filesystem operations.
5. We discuss the digital forensics implications of this new steganographic method.

The rest of this paper is organized as follows: Section (Background) provides a literature review on steganography with emphasis on storage artefacts. Section (Timestamps in modern filesystems) analyzes the use of timestamps on modern filesystems. Section (Steganography based on file timestamps) proposes a novel steganographic channel based on file timestamps. Section (Evaluation of the TOMS system) evaluates the security of the proposed system. Section (Experimental system validation) describes a proof-of-concept implementation aiming at NTFS filesystems. Section (Implications for forensics analysis) discusses the implications on the digital forensics process. Finally, Section (Conclusions and future work) concludes the paper and presents the future directions of our work.

Background

Data hiding

Early works on digital steganography focused on hiding data in the clear, deriving and discussing different methods of embedding data, and arguing how steganography is and probably will be used in the present and in the near future (Katzenbeisser and Petitcolas, 2000; Zielińska et al., 2014). Such works did not anticipate the widespread use of the personal digital devices and the role of the Internet in our daily lives (Franz et al., 1996; Anderson and Petitcolas, 1998).

A considerable amount of research was devoted to embedding unobservable communication within normal network traffic, ranging from the utilization of TCP/IP timestamps (Giffin et al., 2003) to the more general usage of TCP/IP fields (Murdoch and Lewis, 2005). Many implementations of steganography hide encrypted data in innocent-looking network traffic (e.g., ptunnel (Stodtle)), header fields (Rutkowska, 2004), or use timing intervals and artificial transmission delays for information transmission (Lee et al., 2014; Berk et al., 2005; Mazurczyk and Lubacz, 2010). While it has been shown that secure

steganographic protocols are feasible, we are still lacking functional implementations and widespread use of such tools (Hopper et al., 2009).

A second line of research focused on embedding unobservable information within the contents of stored files, introducing undetectable degradation of multimedia quality (e.g., manipulating the low-significance bits of pixel representation in images (Bailey and Curran, 2006)), the color palettes in GIF images (Fridrich and Du, 2000), or (possibly) encoding information in YouTube videos that look like static snow (Williams, 2015).

Filesystems

A plethora of different filesystems is available, including FAT and NTFS for Microsoft-Windows-based devices, ext4 and btrfs for GNU/Linux systems, and HFS+ for Apple OS X and iOS devices.¹ Most of them store different artefacts at various levels of granularity and detail, collectively known as “filesystem metadata”.

Filesystem metadata can be classified in five categories: *file system*, *application*, *file name*, *content*, and *generic* metadata (Carrier, 2005). *File system* metadata are information on how the filesystem is to be read and where the important data structures reside. *Application* metadata are information useful for the application utilizing the filesystem, such as the file owner and the file access permissions. *File name* metadata are information for the human-readable names mapping to logical data locations. *Content* metadata are information about the logical addressing of the files, the file allocation status, and the actual data of the files. *Generic* metadata are information mostly used internally by the filesystem for its operations. This includes information such as the timestamps of various events in the lifecycle of a file.

Steganography using filesystem metadata

The topic of hiding data in filesystem metadata was heavily discussed in the late 1990s (Anderson et al., 1998). Back then, export restrictions on the use of strong cryptographic algorithms outside the USA were in place, and there was an increased concern by the public regarding key escrow. StegFS, a steganographic filesystem compatible with the Linux ext2 filesystem, was developed (McDonald and Kuhn, 2000; Pang et al., 2003). This filesystem achieved plausible deniability of the hidden content thanks to its indistinguishability from unused content. This behavior was achieved by applying encryption on the content under the assumption that a good encryption algorithm ensures that encrypted data appear as random data. However, the use of StegFS is not undetectable as the needed filesystem driver is not hidden. Additionally, there is no integrity check of the data. Thus, StegFS cannot recover from any kind of intrusive data modifications.

Encoding (hiding) information in the order that a filesystem indexes the file names is explored in (Aycock and

¹ An exhaustive list is provided in the Wikipedia entry available at https://en.wikipedia.org/wiki/Comparison_of_file_systems.

de Castro, 2015). The approach is applicable only in the case of a FAT filesystem and cannot be generalized. The file fragmentation is explored in (Khan et al., 2011). This approach introduces significant performance penalties, more evident in magnetic storage media, in the form of delays when accessing the file contents. This delay is due to the heavy file fragmentation that is enforced in order to create the steganographic channel. The delay and the heavy fragmentation can serve as indicators for the presence of steganographic information, thus *defeating* the steganography. Furthermore, (automatic) defragmentation of the storage medium can *destroy* the steganographic information.

Application metadata (e.g., the file owner or the file access permissions) can encode only a few bytes of information and the encoding is easily detected. For example, it is technically feasible to attach an arbitrarily large list of user–permission pairs in an NTFS file, even by referencing non-existent users (Perklin). However, the mismatch of the users mentioned in the system user list and the user–permission pairs, on top of having such long lists in first place, would raise suspicions in a forensics investigator.

The file name cannot be considered as a good candidate for steganographic operations. Indeed, an odd pattern of filenames will look instantly suspicious.

Mixing steganographic information with the actual contents of a file is studied extensively (Katzenbeisser and Petitcolas, 2000; Cheddad et al., 2010; Li et al., 2011; Amirtharajan et al., 2012; Hussain and Hussain, 2013). Format containers for multimedia content (e.g., audio or video) are transparent to and independent of the underlying filesystem that hosts the multimedia file. Thus, a filesystem-level analysis will not be able to disclose the presence of steganographic information in a format container. Also, we note that multimedia transcoding can effectively *destroy* the steganographic information without significantly affecting the original information channel.

Generic metadata, such as temporal information describing the lifecycle of a file, are very sensitive to both the actions of the user and the operating system itself. For example, certain timestamps of file events can be overwritten at any moment while using the filesystem in a normal way. This includes a timestamp of the (last) file modification and (last) access of the file. The fragility of the temporal information might be the reason why, to the best of our knowledge, timestamps have not yet been explored as a steganographic means.

Timestamps in modern filesystems

We analyze in the following paragraphs how modern filesystems use timestamps. The assumption we seek to validate is that there is unused (redundant) capacity in timestamps that is sufficient enough to create a logical channel with steganographic strength.

NTFS is the standard filesystem for Microsoft Windows operating systems. NTFS uses the number of 100 ns passed since January 1, 1601 for its timestamps (Microsoft Developer Network). The timestamps are saved as 64-bit values in the \$Standard_Information field of the Master

File Table (MFT). Additionally, they are saved in the NTFS attribute \$FILE_NAME. Each file has four 64-bit timestamps: (i) creation of the file, (ii) last access of the file, (iii) last modification of the file, and (iv) last modification of the corresponding MFT entry.

ext4 is the successor of the Linux standard filesystem ext3. Ext4 uses 64-bit values to represent timestamps with nanosecond granularity (Mathur et al., 2007; Fairbanks, 2012; Kerrisk, 2010). Ext4 uses the following four timestamps per file: (i) creation of the file, (ii) last modification of the file, (iii) last access of the file, and (iv) the last attribute modification of the file (e.g., permissions or ownership).

btrfs is the upcoming major filesystem for the Linux operating systems (Rodeh et al., 2013). It is a “copy-on-write” filesystem based on B-trees. All file timestamps in btrfs support nanosecond granularity and are saved as 64-bit values (Poeschel and Gim). The first 32 bits of the timestamps are the seconds since the epoch (January 1, 1970) and the remaining 32 bits are the nanoseconds since the beginning of the second. The provided per-file timestamps include: (i) creation, (ii) last modification, (iii) last modification of the file’s attributes (e.g., permissions or ownership), and (iv) last access.

ZFS is intended to be a highly performing, decentralized filesystem (Rodeh and Teperman, 2003). The following per-file timestamps of ZFS have a nanosecond granularity, saved in 64 bits each: (i) creation, (ii) last modification, (iii) last access, and (iv) the last attribute modification (Sun Microsystems Inc., 2006).

FAT32 is the predecessor filesystem of NTFS on the Microsoft Windows operating system. FAT32 uses three different timestamps per file: (i) creation, (ii) last modification, and (iii) last access. The first two timestamps are saved as 32-bit values and the last one is saved as a 16-bit value. The difference comes from the fact that the first two timestamps are provided with a granularity of 2 s, whereas the date of last access is provided with a granularity of one day (Garfinkel, 2012).

HFS+ is the standard filesystem for the Apple Macintosh and iOS devices. HFS+ uses the following per-file timestamps: (i) creation, (ii) content modification, (iii) last attribute modification, (iv) last access, and (v) the last backup (Burghardt and Feldman, 2008). All of these timestamps have a granularity of 1 s and are saved as 32-bit values.

ext3 is the successor of the ext2 filesystem and enhances it by providing journaling capabilities. Ext3 uses three timestamps per file: (i) last access, (ii) last modification, and (iii) last attribute modification. The timestamps have a granularity of 1 s and are saved as 32-bit values. The use of the undocumented large-size inode feature can increase the granularity of the timestamps to 1 ns (Antsilevich, 2009).

Table 1 summarizes our analysis. We confirm that many modern filesystems use 64-bit values as timestamps and offer sub-second granularity (Antsilevich, 2009). This statement covers all filesystems that mainstream consumer operating systems use or access nowadays (e.g., Apple OS X, Google Android, GNU/Linux, and Microsoft Windows) with the exception of the HFS+.

Table 1
Characteristics of filesystem timestamps.

Filesystem	File timestamp	Size	Granularity
NTFS	creation	64 bits	100 ns
	access	64 bits	100 ns
	modification	64 bits	100 ns
	modif. of MFT entry	64 bits	100 ns
ext4	creation	64 bits	1 ns
	access	64 bits	1 ns
	modification	64 bits	1 ns
	attribute modif.	64 bits	1 ns
btrfs	creation	64 bits	1 ns
	access	64 bits	1 ns
	modification	64 bits	1 ns
	attribute modif.	64 bits	1 ns
ZFS	creation	64 bits	1 ns
	access	64 bits	1 ns
	modification	64 bits	1 ns
	attribute modif.	64 bits	1 ns
FAT32	creation	32 bits	2 s
	access	16 bits	1 day
	modification	32 bits	2 s
HFS+	creation	32 bits	1 s
	access	32 bits	1 s
	modification	32 bits	1 s
	attribute modif.	32 bits	1 s
	backup	32 bits	1 s
ext3	access	32 bits	1 s
	modification	32 bits	1 s
	attribute modif.	32 bits	1 s

Three file timestamps, namely *creation*, *access*, and *modification* are supported by almost all the analyzed filesystems. All three timestamps store date and time information with sub-second granularity (one or 100 ns).

The nanosecond precision is not communicated, explicitly or implicitly, to the end users who access the filesystem. They are confronted with file timestamp information that resolves to a second granularity, as depicted in Fig. 1. Thus, there is an information gap between how timestamps are stored and how they are used.




Fig. 1. How file timestamps are displayed to Apple OS X users.

The *creation* timestamp is by and large a static piece of information, as it refers to a unique event, the creation of the file itself. The *access* and *modification* timestamps are updated each time a file is accessed or modified.

Modern operating systems are exploiting latest advances in storage technologies to deliver increased performance and reliability while reducing costs. USB flash drives and SSD storage media are commonplace. In this setting it is advisable, if the application scenario allows so, to reduce filesystem overheads for timestamp housekeeping. This includes disabling the update of per-file access and/or modification timestamps. Such an approach can increase both the performance and the lifetime of a storage medium. Yet, in most consumer-grade usage scenarios, we can expect that only the *access* timestamp remains intact.

The analysis above validates the first part of our initial assumption: there is unused (redundant) capacity in filesystem timestamps. Depending on the filesystem and usage scenario, this capacity ranges from one to nine bytes per file. With modern filesystems hosting hundreds of thousands or even millions of files, this provides enough capacity for storing up to a few megabytes of extra information. In the next sections, we explore how the available capacity can be utilized to create a logical channel with steganographic strength.

Steganography based on file timestamps

We assume a threat model where the attackers can inspect the file contents and can manipulate the filesystem metadata. Also, the attackers can freely remove, rename, or insert new files in the filesystem, and they accept the associated risk of thereby disclosing their presence.

We aim for a steganographic storage system based on file timestamps, namely TOMS (Time-On-My-Side) that is stealthy, robust, and applicable in a wide range of scenarios. “Stealthy” means that the attacker cannot deduce the presence or absence of steganographic information by examining the timestamps. Thus, the attackers are left only with the option of a denial-of-service attack i.e., to overwrite all timestamps and destroy the steganographic channel, thereby disclosing their presence. “Robust” means that the system can sustain and recover from file manipulation attacks. “Widely applicable” means that the system can be configured to match different operation scenarios, balancing performance and secrecy.

System design

The aim of the TOMS system is to hide an input (the *message*) inside the metadata of the hosting filesystem (the *carrier*). For the sake of clarity, we assume that the system can identify the necessary space (i.e., the file timestamps to use) and that all the necessary space is already available. We will return on this issue at the end of the design description (cf. Section (Information representation: the case of NTFS)).

The design of TOMS follows a layered approach. From top to bottom, the system comprises: (i) a storage container layer for the message, (ii) an error correction layer for redundancy, and (iii) an encryption layer.

Storage container layer

The storage container layer maps the message into the underlying file timestamp metadata elementary storage units. The naïve approach for keeping track which files and directories have been used to embed the information is to keep an *encrypted metadata file* with the absolute paths of the files and directories used. The metadata file approach has the benefit that the correct ordering of the files to extract the information is trivial. Also, this file does not necessarily need to be stored in the same filesystem with the hidden data. Rather, it can be stored in another storage media altogether. This is beneficial, since the very presence of a metadata file inside the examined filesystem is neither elegant nor stealthy, even if its contents are encrypted. On the contrary, such an encrypted file can raise further suspicions.

A second option is to reliably embed and extract information *only* based on the files and their timestamps. This can be realized using *oblivious replacements* on whole filesystems (e.g., an NTFS volume) or on the subfolder level (e.g., an NTFS non-root directory). In this case, all files and directories are sorted by their creation timestamp, either globally (filesystem level) or locally (subfolder level). This ordering defines a (logically) continuous storage space that can be used to write and later read the hidden data.

Error correction layer

The normal use of the storage medium hosting the filesystem as well as the actions of the attackers may remove some of the files stored on the filesystem. Also, the attackers might intentionally change the creation timestamps of some of the files. Such actions, deliberate or not, cause a new ordering of the creation timestamps, which results in the inability to either access the input file segments in the correct order or altogether.

The error correction layer augments the initial representation of the input file with additional information that can cope with the aforementioned issues. As a first step, an error correcting code (ECC) is appended to the representation. The ECC can both detect and reconstruct missing information. As a second step, this layer enforces a *start offset* for the used files. This allows the program to start from a random point in the ordering and use both older and newer files. Thus, not only old files are used to hide information.

The selection of an appropriate ECC is left to the implementation. By and large, an ECC should not introduce significant storage overhead.

Encryption layer

The error correction layer introduces data redundancy. This redundancy comes on top of the structured information needed to represent the links from timestamp to timestamp in order to form a logically continuous storage space. These can be sources that result in timestamps with patterns. If patterns are detected, the whole steganographic system will collapse, since they reveal the existence of hidden information.

The role of the encryption layer is twofold: On the one hand it protects the hidden information from disclosure. Only somebody in possession of the related cryptographic

key(s) can access the encrypted and hidden information. On the other hand the confusion and diffusion properties of a (good) secure cipher ensure that hardly any pattern will exist in the output allowing it to appear totally random.

The encryption layer uses symmetric-key cryptographic algorithms to encrypt the information of the two previous layers. Stream ciphers, as for example AES-OFB or RC4, can be used in this layer. The advantage of stream ciphers over block ciphers is that the former do not need to expand the output of the operation and that they can recover to a certain point from errors (e.g., missing timestamps) at a bit rather than a block (i.e., dozens of bytes) level.

Information representation: the case of NTFS

We can now describe how the TOMS components work together to hide a message in the file timestamps. In the following, we will use the NTFS filesystem as an example. However, the description is valid for any other filesystem with similar characteristics. Fig. 2 depicts the NTFS inode data structure used to represent various filesystem objects, including a file and a directory. In the following, we will use the term “file” to refer to NTFS inodes.

Two file timestamps can be used by TOMS in the case of NTFS: the *creation* and the *last access*. Each timestamp uses 24 bits to represent its nanoseconds part. Thus, a total of six bytes per file can be used to hide information. This constitutes the elementary storage unit (ESU) for the TOMS system. We assume that the size of the input (steganographic) message, M , is much larger than the size of an ESU. First, an error correcting code function is applied to the input message, $E = ECC(M)$. Then n , the number of ESUs needed to store E , is prepended ($n||E$).

Information hiding

The information hiding process works as follows. The encoded message (E) is fragmented into n blocks of five bytes each (B_1, B_2, \dots, B_n). Then, every block is prepended with one byte that is used as a block index ($i \in \{1, \dots, n\}$). The special value of 0x00 for the index byte is used to prepend a block of five bytes that contains the number of needed ESUs, n . The resulting structure is a linked list of six-byte blocks: $(0, n), (1, B_1), \dots, (n, B_n)$. This structure is then encrypted with a stream cipher and a secret key, producing an output list of six-byte blocks: C_0, C_1, \dots, C_n .

TOMS constructs the list of candidate files that they can be used as carriers. The list, F , is ordered based on the creation timestamp of each file, and a start offset, s , is chosen randomly. The ordered list of files, $F_s = \{f_s, f_{s+1}, \dots, f_{s+n}\} \in F$, will be used as the carrier. TOMS then proceeds and replaces the nanoseconds part of the creation and access timestamps of each file in F_s with the six-byte encrypted chunk C_i .

Large message handling

Using just one byte as index limits the length of the hidden message (E) to only 255 bytes. We overcome this limitation by allowing multiple index bytes to share the same value (overflow upon reaching the value 0xFF and restart numbering from 0x01). Whenever an overflow occurs, an ESU is consumed in order to store the length of the




Fig. 2. Overview of storing data in the nanoseconds part of the timestamp fields.

whole message, using again an index byte of 0x00. Thus, every ESU with an index byte holding the value 0x00 contains the *total* length of the message.

Recall of hidden information

The information recall process works as follows. The timestamps for *all* the files in the filesystem are extracted, sorted by their creation time, and then saved as a list *G*. For every list entry, the nanoseconds part of the creation and the access timestamps are decrypted by applying the same stream cipher and key material used during the information hiding process. If the decrypted first byte of the creation timestamp equals the index byte value 0x00, the respective timestamps are added in an offset list *L* and the number of ESUs, *n*, is recovered. Then, the next *n* list entries are processed, recovering the respective index (i.e., 0x01, 0x02, ..., *n*) and the structure *H*. Next, the error correction code function is applied on *H*, recovering the original (hidden) message *M*.

Evaluation of the TOMS system

In this section, we evaluate the design principles of the TOMS system. Our evaluation is based on theoretical, experimental, and evidence-based analysis of the steganographic strength of TOMS under the assumed threat model (cf. Section (Steganography based on file timestamps)).

Stealthiness

“Stealthiness” describes the degree to which the very existence of hidden information is disguised, irrespectively of the ability to recover the hidden message(s). We analyze the two factors defining the stealthiness of the TOMS system in the following.

Timestamp handling by operating systems

The use of the encryption layer ensures that the steganographic information are not recoverable without having access to the key material of the stream cipher used.

Furthermore, a good stream cipher ensures that each output bit will be a one or a zero with equal probability. But how do modern operating systems handle the timestamp information in first place? Do they fill these data structures with sub-second-precise information or do they opt for a different strategy? If the former, what is the precision of the provided time information? We sought the answer to these questions using three approaches.

Code audit

As a first approach, we performed a code audit of the NTFS-3g implementation of the NTFS filesystem (Tuxera Inc). This is the default driver for accessing NTFS volumes from within the Linux and Apple OS X operating systems, and it is an open source code. A similar code audit for the NTFS implementation of the Microsoft Windows operating system was beyond our reach, since the source code is not publicly available. The code audit revealed that the NTFS-3g fills the related timestamp structures with nanosecond-granular information provided by the Linux system clock which also has a nanosecond granularity.

Synthetic data

The second approach was to create synthetic data for experimentation, which is online at the authors' website together with the most recent version of the code used for this purpose.² We created files in batches using a Python script on a Linux system accessing an NTFS volume via NTFS-3g. Each batch created 100,000, one million, or ten million different files. Half of the files were created with a random delay of one to 2 s between each creation. The other half of the files were created with zero delay, i.e., as fast as the computer system could handle the requests.

Our “synthetic” dataset contains 117 million files. We collected the three timestamps (create; access; and modify, all equal to each other) for this dataset as well. We conducted an exploratory data analysis to determine if the timestamp distribution was uniformly distributed. Results

² <https://www.sba-research.org/dfrws2016/>.

for the Kolmogorov–Smirnov goodness-of-fit test for uniformity indicated that the timestamp distribution did not deviate significantly from a uniform distribution ($D = .99, p = 2.2 \times 10^{-16}$).

The source code audit and the experiments validate that the Linux operating system uniformly uses the full spectrum of the 24-bit sub-second granularity to store the timestamp information.

Real data

The third approach was to collect evidence from Microsoft-Windows-based systems that are actively used to perform day-to-day tasks (“real-world systems”). We therefore collected the file timestamp information from a sample of 70 filesystems (NTFS volumes) from multiple Microsoft Windows computers available at our research lab. On average, each filesystem of our sample contained about 290,000 files and 40,000 directories; the largest one contained over 2.2 million files and directories. The majority of the sampled filesystems ($n = 63$) were actually “system volumes”, i.e., they contained the files of the Microsoft Windows operating system (e.g., those files commonly found in the C: \ Windows directory) and (most likely) of the majority of installed software (e.g., those files commonly found in the C: \ Program Files directory). Only seven systems contained more than one NTFS volume (i.e., “non-system volumes”). Such volumes are often used as storage for work or personal data (e.g., documents, spreadsheets, and pictures). In total, our “real-world” dataset contains the timestamps of 22,261,386 files and directories.

We analyzed the timestamps contained in this dataset and we noticed some irregularities in their distribution. Creation timestamps that are filled with zeroes in their nanoseconds parts were disproportionately more frequent than expected. This is the case when files are migrated into NTFS volumes from FAT32 filesystems. The latter use a granularity of 2 s at best, hence the zeroes. This assumption was empirically tested and further confirmed by Microsoft’s documentation regarding timestamp changes (Microsoft).

Time of filesystem inspection

In the previous paragraphs, we saw that the timestamps can be used as stealthy information carriers, since the sub-second information follows a uniform distribution, as does the output of a stream cipher encryption. Before replacing any timestamps, one should consider if and how often the filesystem is inspected by an attacker. As an example, we consider the case of (operating) system files. These files are installed once and are seldom, if ever, touched again (e.g., only when operating system updates are installed). Thus, if their timestamps are proactively collected, any future modification by the TOMS system will be easily detected.

In a forensics analysis scenario, we can assume that the investigator will inspect the metadata after the message was hidden in the timestamps. We can also assume that the investigator does not have access to earlier versions of the filesystem metadata information. Thus, existing timestamps can be utilized to hide steganographic information. In a scenario where the filesystem can be proactively

inspected, already existing files might not be good candidates for carriers. Thus, only new files (i.e., generated after the last inspection) can be utilized to hide steganographic information.

We assume for our subsequent analysis a more conservative scenario, in which the filesystem is *proactively* inspected. In this case, one should exclude all system files and all files with timestamps containing zeroes in the sub-second part. Applying this to our initial real-world dataset, it resulted in an almost 50% drop of available files, down to 11 million files spanning 70 NTFS volumes.

Robustness

“Robustness” refers to the ability of the TOMS system to cope with changes in the filesystem contents. The information hiding and recall procedure of TOMS is straightforward when the initial ordered list of files F_s remains intact between information hiding and information recall(s). In the following, we analyze how the TOMS system defends against actions that result in modifications of F_s .

Encrypted metadata

This is the simplest of the proposed storage container layers. The ordered list of files F_s is not affected by operations on the filesystem level (assuming that these operations do not touch the timestamps). Should some files have been removed from the filesystem, or some timestamps are updated, the encryption and the error correction layers may be able to recover the lost information thanks to the use of the stream cipher and the ECC. If and how the recovered information is stored back to the timestamps (e.g., insert new files, re-encode information, or even fix the “corrupted” timestamps) is a decision to be made taking into account the severity of the errors and the assumed time and frequency of inspection.

Oblivious replacements

In this approach the ordered file list F_s results from sorting the timestamps that are provided by the filesystem. Thus, it might be the case that the TOMS system unknowingly uses a different list F'_s for information recall instead of the one originally used for information hiding. If some files were removed between information hiding and recall(s), the same arguments as in Section (Encrypted metadata) apply.

We assume that some additional files, F_a , are included in the $F'_s = F_s \cup F_a$ list, and that the computing system has a proper clock. If oblivious replacement is applied globally (filesystem level), the TOMS system will always recover the correct F_s . This is feasible because the file creation timestamp is immutable on an NTFS volume, i.e., it does not change when the file is copied, renamed, or moved *within* the same NTFS volume (Carrier, 2005). Thus, even if files are moved across different NTFS folders, their creation timestamp will not change. Also, these additional files F_a will have more recent creation timestamps than those already contained in the original F_s and allow therefore a clear separation of the two sets.

If oblivious replacement is applied locally (subfolder level), then it is possible that the ordering of F_a is

intermixed with the ordering of F_s . This is the case where files are moved to the specific subfolder containing F_s ; as mentioned earlier, the file creation timestamps are immutable. This situation is accommodated by the use of the encryption and the error correction layers. Assume that a file $f_m \in F_a$ is inserted in the ordered list F_s . First, the ESUs of f_m must decrypt correctly and not be rejected by the encryption layer. Then, the value of the index byte contained in the (erroneously) decrypted ESUs must match the currently expected sequence number in order to not be rejected by the error correction layer. Finally, the payload information contained in the ESUs must pass through the ECC. Only then, these information are accepted as valid. If the processing of f_m fails, the error correction layer provides the necessary protection to recover from the error. Thus, the two layers provide an adequate defense (up to a certain point) against such (deliberate or not) insertion attacks. The amount of redundant information handled by the ECC defines this protection level. An oblivious replacement at the subfolder level requires a stronger ECC compared to the filesystem level.

Applicability

“Applicability” refers to the degree to which the TOMS system can be utilized in various application scenarios. The layered design of TOMS provides an initial indication of its wide applicability. The TOMS system supports three different storage layers and is agnostic to the ECC used as well as to the stream cipher. Furthermore, TOMS can be easily applied to any modern filesystem that supports sub-second timestamp granularity; while the basic description supports two timestamps often found in modern filesystems (namely, *creation* and *last access*), there is no design constraint regarding the number of timestamps or the use of filesystem-specific timestamps (e.g., *last attribute modification* for ext4). The design of the TOMS system allows to explore various performance tradeoffs in order to match the secrecy requirements of the selected application. We discuss these tradeoffs in the following paragraphs.

The application scenario defines the use of existing files or opts to create new ones to embed a steganographic message. In the latter case, it is advisable to generate small files that act as carriers (e.g., files in the range of few thousands bytes).

The use of an ECC introduces an overhead of 10–20%. If the risk of information loss can be sustained, the use of an ECC can be omitted altogether.

The selection of the storage container type is important. If an encrypted metadata file is used, one must decide if the contents of the file should be embedded in the filesystem or stored elsewhere. The resulting size of this metadata file can be a decisive factor. When embedding about 1.5 MB of data into 175,000 timestamps, the corresponding metadata file takes about 215 KB of disk space. A benefit of this approach is that there is no need to store index bytes to rebuild the ordered list of carrier files and recover the hidden information. Also, file reordering is not a threat in this case (unless someone is tampering with the metadata file) and thus the performance requirements for an ECC are more relaxed (or can be omitted altogether).

The oblivious replacement approach mandates the use of index bytes. Each ESU uses one index byte per five payload bytes (ratio of 1:5). If only one timestamp is available for each file, the ratio becomes 1:2, which may cause a lot of overhead. On the other hand, if three timestamps are available, this ratio becomes 1:8, which is quite efficient. Compared to an encrypted metadata approach, oblivious replacement needs between 12.5% (two timestamps) and 20% (three timestamps) more files in order to store the same amount of hidden information.

Experimental system validation

We developed a proof-of-concept implementation of the TOMS system for the experimental validation of our steganographic proposal. The implementation targets the NTFS filesystem and is based on the Python language version 2.7 for flexibility and increased portability. Our implementation can be delivered as a stand-alone executable and does not require the installation of special software or any modifications of the Linux kernel. It realizes the layered design described in Section (Steganography based on file timestamps) and can be easily ported to work with any filesystem that uses a nanosecond timestamp granularity.

The development and experimentation platforms are based on Xubuntu Linux 15.04 64-bit, running kernel version 3.19.0–25, the latest stable one at the time of writing. The underlying disk on which the operating system is installed is a solid state disk (SSD) for faster I/O access. As NTFS is Microsoft-proprietary, we opt for NTFS-3g in its current version. The steganographic executable application takes care of all information management tasks. The application is assumed to have full access to the NTFS volume (filesystem).

The application supports the use of two- and three-file timestamps. The file *creation* and *last access* timestamps are not modified by the operating system: starting with Microsoft Windows Vista, the default value of `NtfsDisableLastAccessUpdate` is set to one (Hermann). The corresponding mount option in Linux is `noatime`; in most of the popular Linux distributions this option is not activated by default. The file *last modified* may be modified under normal use, so it is up to the users to decide if they enable it (and pay attention not to destroy the related information during the normal use of the filesystem).

Information hiding and recall

The typical work flow for information hiding is as follows: the user starts the Python application and provides (i) the message to be hidden, (ii) a key to encrypt the message, (iii) the method for hiding (metadata file, oblivious replacements on volumes, oblivious replacements on subfolders), and (iv) the number of different timestamps to use (either two: *creation* and *access* or three: *creation*, *access*, and *modification*). Once the necessary information are collected, the application performs the following steps: (i) it concatenates the message with the error correction code, (ii) adds the index bytes to the resulting data (if the chosen hiding method is not the metadata file), (iii) encrypts the

data with the stream cipher, and (iv) embeds the encrypted data into the timestamps. On the information recall path, the user enters the encryption key and the application displays the decrypted message.

Listing 1 and Listing 2 outline in pseudocode the two work flows for information hiding and recall respectively.

Listing 1: Embedding data in timestamps.

```

1 def hide(path, msg, key):
2     rs= calcReedSolomon(msg)
3     m= msg · rs
4     C= chunk(m, 5)
5     index= 0
6     temp= ∅
7     for c ∈ C:
8         s= c
9         if index = 0 or index % 255 = 0:
10            s= length(m)
11            temp= temp · index · s
12            index++
13     files=sort(recEnumFiles(path), by=creation_time)
14     offset= calcRandomOffset()
15     while offset:
16         files.pop()
17     em= encrypt(temp, key, mode=RC4)
18     C= chunk(em, 6)
19     for c ∈ C:
20         f= files.pop()
21         f.creation_time.nsec= c[0:3]
22         f.access_time.nsec= c[3:6]

```

Listing 2: Extracting data from timestamps.

```

1 def extract(path, key):
2     F= sort(recEnumFiles(path), by=creation_time)
3     em= ∅
4     for f ∈ F:
5         c= f.creation_time.nsec · f.access_time.nsec
6         em= em · c
7     m= decrypt(em, key, mode=RC4)
8     C= chunk(m, 6)
9     i, l= 0, 0
10    for (i, c) ∈ enumerate(C):
11        if c[0] ≠ 0x00:
12            continue
13        l= int(c[1:6])
14        break
15    S= sort(C[i:l], by=first_byte)
16    temp= ∅
17    for (i, c) ∈ enumerate(S):
18        t= c[1:6]
19        if c[0] ≠ i:
20            t= 0x00 × 5
21        temp= temp · t
22    return decodeReedSolomon(temp)

```

Metadata file information protection

All information processed by the application are held in memory (RAM) and are encrypted with AES-256-CBC using a user-provided key. This is a precautious measure in order to protect against extraction of the plain metadata file during a forensics analysis of the storage medium, e.g., in the slackspace of the hard disk (Mulazzani et al., 2013).

After the information has been embedded, the metadata file is built from the information kept in RAM so far. Before writing this information to the disk, it is compressed using gzip and encrypted with the AES algorithm using a user-provided password. Our implementation supports the use

of different passwords for the metadata file and the actual data.

We take care not to accidentally write the unencrypted metadata file to the disk, as this could leave persistent traces which particular files were modified. During the embedding process the information resides unencrypted in the RAM, and we did not implement countermeasures to prevent the operating system to store the corresponding memory pages on the disk, e.g., due to paging or hibernation. However, our application supports the encryption of information on a per-path basis, right after embedding the information in order to minimize the time the unencrypted information resides in the RAM, at the cost of creating a much larger metadata file due to the lack of compression.

Performance

Two of the main considerations of steganographic systems are the undetectability and the confidentiality of the hidden data (Morkel et al., 2005). The performance of the system is also an important factor with respect to applicability.

We performed a series of experiments to gain insights regarding the performance of TOMS when embedding and extracting information using volume-wide oblivious replacement. Table 2 summarizes our findings. The reported figures are the averages of ten consecutive executions of hiding (embedding) and recall (extracting). The amount of space used to embed data is reported as a percentage of the overall available storage space provided by the ESUs (i.e., 6 bytes per file). The time needed to hide (embed) the information is almost constant, irrespective of the data volume. On the other hand the time to recall (extract) the information is almost linear to the percentage of embedded data. In both cases the time ranges in dozens of seconds, which might be considered too high. Upon closer inspection, it appears that the calculation of the Reed-Solomon ECC dominates the processing time for both. However, since the file metadata are extracted from the MFT, which resides in the RAM, the average time to extract is lower than the average time to embed. This lower time is caused by performing the vast majority of filesystem operations within the RAM instead of directly accessing the hard disk.

Effect on actual filesystem operation

As a final consideration, we examined if the filesystem remained operational for normal use after manipulating the stored file timestamps. We mounted and unmounted the NTFS volumes that were modified by our proof-of-concept implementation using the drivers provided by

Table 2

Time to embed and extract information on filesystem.

Space used	15%	30%	50%
Timestamps needed	78,687	157,325	264,193
Time to embed [sec]	74.78	76.17	76.33
Time to extract [sec]	20.19	36.92	60.29

Linux, Microsoft Windows, and Apple OS X operating systems. We did not notice any problems in using the volumes, and no error messages were logged by the operating systems. We also performed regular file operations in the volumes and did not notice any issues. Recall of the steganographic information after the regular use succeeded without any problems as well.

The analysis validates our initial assumption: typical usage scenarios of modern filesystems allow to persistently store additional information in file timestamps without affecting their normal use.

Implications for forensics analysis

Responsible research in steganography involves both developing new techniques for information hiding and detection (steganalysis). The issue of detection is increasingly important for digital forensics examiners, as criminal activities through digital means are becoming prevalent (Europol, 2015).

Embedding information in file timestamps is feasible. As discussed in Section (Evaluation of the TOMS system), these information are indistinguishable from that of normal operations, provided that a stream cipher is used in the encryption layer. Thus, a statistical analysis of file timestamps should be incorporated in the forensics examination procedures as a first line of defense. This analysis can provide hints for the presence of information hidden in the timestamps, if one opts to disable the encryption layer of the TOMS system.

There may be additional artefacts and implementation details which can assist a forensics investigator in disclosing the presence of hidden information. A fully-functional TOMS demands careful implementation and operation decisions. The developer must ensure that the application leaves no installation or execution traces that could reveal its presence. If a backup image of the filesystem contents at an earlier time is available to the examiners, they can compare the timestamps regarding unjustifiable modifications, especially for the case of the creation timestamps. Furthermore, the generation of new files to use them as carriers must be justifiable from a modus operandi point of view. For example, if modification timestamps are utilized, it must be justifiable why they differ from the creation timestamps – this would be suspicious if it occurs in the same second for a large batch of files. It is advantageous to check for modus operandi violations during the forensics examination process.

Another approach for the investigation procedure is the correlation of an installation timeline for an operating system and its well-known application files (e.g., the Microsoft Office suite). If such files are used to hide information and if they share the same timestamp up to the second part, an installation timeline can reveal that the creation order of some files does not match the expected one.

The filesystem data structures are not the only place where timestamps are stored. If an operating system records file-related events in its system logs with nanosecond precision, a digital forensics investigator can perform a correlation analysis between these two information

sources in order to detect unjustifiable mismatches. Operating systems also use transaction logs (journals) for recovery processes (e.g., the NTFS Transaction Log \$LogFile (Cho, 2013)). These can also be used for correlation analysis. Wiping out such log files or carelessly modifying them can raise further suspicions and assist aiding the investigation along.

In the case of NTFS, an informed decision in our proof-of-concept (PoC) implementation was to modify *only* the filename attribute of the MFT. However, the same information are maintained in the standard information attribute as well. Thus, an investigator can compare the two attributes and detect the use of the PoC implementation.

Conclusions and future work

In this paper, we proposed and explored the applicability of file timestamps as a steganographic channel. Based on our analysis of how modern operating systems store timestamps for file events in filesystem data structures and how they are displayed to the users, we reveal a redundant space to hide information. We described how this space can be utilized as a steganographic channel using a layered design that offers stealthiness, robustness, and wide applicability. We evaluated our design through theoretical, evidence-based, and experimental analysis in the case of the NTFS filesystem with datasets containing millions of files: the hidden information are statistically indistinguishable from timestamps produced during normal use. We also validated the applicability of our proposal through a proof-of-concept implementation targeting the NTFS filesystem. Finally, we discussed the implications of this new steganographic technique for digital forensics analysis.

As future work, we consider it interesting to confirm our findings with datasets provided by other researchers and practitioners in the field. Regarding our evaluation on the NTFS filesystem, we plan to explore implications of the TOMS system on the NTFS \$LogFile. Also, to extend the validation to other filesystems, such as the ext4 that is used natively by the Linux operating system but also showing the robustness of TOMS during heavy usage of the underlying filesystem. Another direction is to explore additional filesystem artefacts and data structures beyond file timestamps that exhibit similar characteristics regarding time handling. It is also useful to extend the study towards devices that provide timing information with smaller precision: such devices may not offer the timing granularity that results in uniformly-distributed timestamps.

Acknowledgments

The research was supported by the Austrian Research Promotion Agency (FFG) through the Bridge Early Stage grant P846070 (SpeedFor) and the COMET K1 program.

References

- M. Akgül, M. Kırıldoğ, Internet censorship in Turkey, *Internet Policy Rev* 4(2).
- Amirtharajan R, Qin J, Rayappan JBB. Random image steganography and steganalysis: present status and future directions. *Inf Technol J* 2012; 11(5):566–76.

- Anderson D. Splinternet behind the great firewall of China. *Queue* 2012; 10(11):40–40:49.
- Anderson RJ, Petitcolas FA. On the limits of steganography. *IEEE J Sel Areas Commun* 1998;16(4):474–81.
- Anderson R, Needham R, Shamir A. The steganographic file system. In: *Information hiding*, Springer; 1998. p. 73–82.
- Antsilevich E. Capturing timestamp precision for digital forensics, Tech. Rep. JMU-INFOSEC-TR-2009-002. USA: Department of Computer Science, James Madison University; 2009.
- Aycock J, de Castro DMN. Permutation steganography in FAT filesystems. In: *Transactions on data hiding and multimedia security X*, Springer; 2015. p. 92–105.
- Bailey K, Curran K. An evaluation of image based steganography methods. *Multimed Tools Appl* 2006;30(1):55–88.
- Berk V, Giani A, Cybenko G. Detection of covert channel encoding in network packet delays, Tech. Rep. TR2005-536. USA: Department of Computer Science, Dartmouth College; 2005.
- Burghardt A, Feldman AJ. Using the HFS+ journal for deleted file recovery. *Digit Investig* 2008;5:S76–82.
- Carrier B. *File system forensic analysis*. Addison-Wesley Professional; 2005.
- Cheddad A, Condell J, Curran K, Mc Kevitt P. Digital image steganography: survey and analysis of current methods. *Signal Process* 2010;90(3): 727–52.
- Cho G-S. A computer forensic method for detecting timestamp forgery in NTFS. *Comput Secur* 2013;34:36–46.
- Cox I, Miller M, Bloom J, Fridrich J, Kalker T. *Digital watermarking and steganography*. Morgan Kaufmann; 2007.
- Europol. 2015 Internet Organized Crime Threat Assessment (IOCTA), online at <https://www.europol.europa.eu/content/internet-organised-crime-threat-assessment-iocta-2015>.
- Fairbanks KD. An analysis of ext4 for digital forensics. *Digit Investig* 2012; 9:S118–30.
- Franz E, Jerichow A, Möller S, Pfitzmann A, Stierand I. Computer based steganography: how it works and why therefore any restrictions on cryptography are nonsense, at best. In: *Information hiding*, Springer; 1996. p. 7–21.
- Fridrich J, Du R. Secure steganographic methods for palette images. In: *Information hiding*, Springer; 2000. p. 47–60.
- Garfinkel S. Digital forensics XML and the DFXML toolset. *Digit Investig* 2012;8(3):161–74.
- Giffin J, Greenstadt R, Litwack P, Tibbetts R. Covert messaging through TCP timestamps. In: *Privacy enhancing technologies*, Springer; 2003. p. 189–93.
- Greenwald G. No place to hide: Edward Snowden, the NSA, and the US surveillance state. Macmillan; 2014.
- U. Hermann, Access Time Update, online at <http://forensicswiki.org/wiki/MAC-times\#Access-Time-Update>.
- Hopper N, von Ahn L, Langford J. Provably secure steganography. *IEEE Trans Comput* 2009;58(5):662–76.
- Hussain M, Hussain M. A survey of image steganography techniques. *Int J Adv Sci Technol* 2013;54:113–24.
- Tuxera Inc., Open source: Ntfs-3g, online at <http://www.tuxera.com/community/open-source-ntfs-3g/>.
- Katzenbeisser S, Petitcolas F. *Information hiding techniques for steganography and digital watermarking*. Artech House; 2000.
- Kerrisk M. *The Linux programming interface*. No Starch Press; 2010.
- Khan H, Javed M, Khayam SA, Mirza F. Designing a cluster-based covert channel to evade disk investigation and forensics. *Comput Secur* 2011;30(1):35–49.
- Lee KS, Wang H, Weatherspoon H. PHY covert channels: can you see the idles?. In: *Proceedings of the 11th USENIX Symposium on networked systems design and implementation (NSDI 14)*, USENIX; 2014. p. 173–85.
- Li B, He J, Huang J, Shi YQ. A survey on image steganography and steganalysis. *J Inf Hiding Multimed Signal Process* 2011;2(2):142–72.
- Mathur A, Cao M, Bhattacharya S, Dilger A, Tomas A, Vivier L. The new ext4 filesystem: current status and future plans. online at. 2007. <https://www.kernel.org/doc/mirror/ols2007v2.pdf>.
- Mazurczyk W, Lubacz J. LACK – a VoIP steganographic method. *Telecommun Syst* 2010;45(2–3):153–63.
- McDonald AD, Kuhn MG. StegFS: a steganographic file system for Linux. In: *Information hiding*, Springer; 2000. p. 463–77.
- Microsoft, Description of NTFS date and time stamps for files and folders, online at <https://support.microsoft.com/en-us/kb/299648>.
- Microsoft Developer Network, File times, online at <https://msdn.microsoft.com/en-us/library/ms724290/>.
- Morkel T, Eloff JH, Olivier MS. An overview of image steganography. In: *The fifth annual information security South Africa (ISSA 2005)*, electronically; 2005. p. 1–12.
- Mulazzani M, Neuner S, Kieseberg P, Huber M, Schrittwieser S, Weippl E. Quantifying windows file slack size and stability. In: *Advances in digital forensics IX*, Springer Berlin Heidelberg; 2013. p. 183–93.
- Murdoch S, Lewis S. Embedding covert channels into TCP/IP. In: *Information hiding*, Springer; 2005. p. 247–61.
- Pang H, Tan K-L, Zhou X. StegFS: a steganographic file system. In: *Proceedings of the 19th International conference on data engineering, IEEE*; 2003. p. 657–67.
- M. Perklin, ACL steganography: permissions to hide your porn, online at <https://www.defcon.org/images/defcon-21/dc-21-presentations/Perklin/DEFCON-21-Perklin-ACL-Steganography-Updated.pdf>.
- S. Poeschel, J.-H. Gim, Btrfs on-disk format, online at <https://btrfs.wiki.kernel.org/index.php/On-disk-Format\#Basic-Structures>.
- Rodeh O, Teperman A. ZFS – a scalable distributed file system using object disks. In: *Proceedings of the 20th IEEE/11th NASA Goddard conference on mass storage systems and technologies (MSST 2003)*, IEEE; 2003. p. 207–18.
- Rodeh O, Bacik J, Mason C. BTRFS: the linux B-tree filesystem. *ACM Trans Storage (TOS)* 2013;9(3):9.
- Rutkowska J. The implementation of passive covert channels in the Linux kernel. In: *Chaos communication congress, chaos computer club eV*; 2004.
- D. Stodler, Ping tunnel, online at <http://www.cs.uit.no/~daniels/PingTunnel/>.
- Sun Microsystems Inc. ZFS on-disk specification. 2006.
- A. Williams, Transfer data via YouTube, online at <https://hackaday.com/2015/08/23/transfer-data-via-youtube/>.
- Zielińska E, Mazurczyk W, Szczypiorski K. Trends in steganography. *Commun ACM* 2014;57(3):86–95.