The Privacy of Private Browsing

Ashley Hedberg[1]

ashley.hedberg@tufts.edu

[1] *Mentor:* Ming Chow, Tufts University

Abstract

Most modern web browsers have a “private browsing” mode that supposedly allows a user to surf the internet without leaving any traces of his or her activity on his or her machine. However, the notion of “private browsing” offers users a false sense of security, as browsing information is often left behind when a private browsing session terminates. Several researchers have already demonstrated methods of detecting this information. These include the analysis of the hard drive, memory, and the pagefile on Windows machines. The existence of private browsing artifacts on a user’s machine raises many questions: how much privacy do web browser developers actually aim to provide with “private browsing” modes? Are these goals accurately conveyed to the end user? And, if the users browsing session can be reconstructed from these artifacts, how much information would digital forensics professionals be able to gain? This paper seeks to answer these questions and to prove that such artifacts can be used to determine, at least in part, what a user was doing during his or her “private browsing” session—thereby rendering it not very private at all.

1 Introduction

1.1 Browser Overview

Currently, the most popular browsers for the Windows family of operating systems are Mozilla Firefox, Google Chrome, and Microsoft Internet Explorer. All three of these browsers have built-in private browsing modes. While they all have different names—"Private Browsing" in Firefox, "Incognito Mode" in Chrome, and "InPrivate Browsing" in Internet Explorer—they all claim to do essentially the same thing: allow a user to browse the web without traces of his or her browsing history, login information, form and search bar entries, or downloads remaining on his or her machine.[8][2][3]

1.2 Definitions

Many shortcomings of private browsing modes relate to how a computer's operating system manages memory and disk space. Those who are unfamiliar with these tasks may benefit from the definitions below.

- *Cluster*: A group of disk sectors that the operating system treats as a unit; smallest unit of disk storage that the operating system can allocate
- *Digital certificate*: Verifies that a particular website or target server is who it claims to be when connecting securely
- *Disk*: The hard drive of a computer; where a user's files and directories (among other things) are stored
- *Free space*: A cluster that is not allocated to any particular file
- *Page*: The smallest unit of memory that the operating system can allocate; can be temporarily stored on disk to allow the computer to run more programs at the same time
- *Slack space*: The space between the end of a file and the end of the cluster in which it is stored; important in computer forensics for finding fragments of old files[1]

[1] Hard drive clusters are of a fixed size determined by the operating system. Most of the time, when a user saves a file to his or her hard drive, the file is not exactly the same size as the cluster in which the operating system puts it. Furthermore, when a user "deletes" a file from his or her hard drive, the file is not

2 To the Community

There are many misconceptions about private browsing. Some believe that it prevents an internet service provider, network administrator, or attackers engaged in packet sniffing from linking a user's internet activities to his or her identity. It does not. Any and all network packets leaving the user's machine contain information such as the user's IP address, which can be used to determine his or her location and/or identity. Others think that private browsing sessions leave absolutely no trace of their activities on their local machines. As the next section of this paper will describe, this is also false. Still others think that private browsing will prevent the National Security Agency from tracking them on social media. It certainly does not. This paper will shed some light on what private browsing sessions can and cannot do, allowing internet users to think twice before exploring the web.

3 Applications: Forensics

One of the biggest problems with private browsing sessions is that artifacts remain on a machine after the user exits his or her internet browser. These artifacts can include logon information for websites, browsing history, and digital media such as images and videos—all items that you would expect a private browsing session to keep private. Additionally, browser extensions can often record this kind of information, even when a user is browsing privately.

3.1 Where To Look

Focused analysis of the hard drive can reveal the most artifacts. One private browsing experiment in which logon information, browser history, and cached images were recovered obtained most of the artifacts from the hard drive's free space and slack space.[4] The

really deleted—it is simply marked as available to be overwritten if necessary. If a file is eventually stored on the hard drive in a cluster that used to be home to a larger (now "deleted") file, leftover data from the old file will still exist at the end of that cluster.

Windows pagefile (pagefile.sys on the hard drive) in particular can reveal private browsing artifacts. This file is where the least recently used pages of memory are stored when too many applications are competing for the computer's physical memory. Browsing history and keywords used in internet search engines have been discovered here.[5] Free programs such as Autopsy allow anyone to investigate the free and slack space of a hard drive.

Artifacts can also be recovered from the computer's memory itself. Cached web documents can often be found here.[4] Browsing history, logon information, and cookies have been detected after private browsing sessions by analyzing a computer's memory. The SQLite databases used by Chrome and Firefox during browsing sessions are generally only stored in memory when private browsing is in use, and their residual data could be detected, as well.[6]

3.2 Mozilla Firefox

In 2011, researchers discovered that Firefox 3.6.11 stored browser history and search engine keywords in the computer's physical memory and that this information could be accessed after the browsing session was terminated. An analysis of the contents of pagefile.sys revealed similar information. [5] The presence of these in-memory artifacts persisted through Firefox 17.0.1.[4] Additionally, digital certificates received during private browsing sessions in Firefox 3.6 were not removed from their on-disk database once the browsing session was terminated.[1]

Researchers demonstrated in 2013 that on-disk files used by Firefox 17.0.1 showed modified timestamps after a private browsing session. Even though information about the user's activities may not be directly accessible, evidence that a private browsing session took place can help forensic investigators determine why browsing history cannot easily be located on a suspect's machine, especially when analysis of the computer's memory is not possible.[4] Firefox 19.0 fixed this problem.[6]

3.3 Google Chrome

One 2011 study stated that it was more difficult to retrieve private browsing artifacts from Chrome's incognito mode (using Chrome 7.0.517.41) than Firefox or Internet Explorer. That study found that, as in Firefox 3.6.11, browsing history and search engine keywords could be accessed in the computer's memory after the incognito session terminated. However, it found no on-disk artifacts. Only physical memory contained these artifacts, as the pagefile revealed no traces of the websites visited or searches performed using incognito mode.[5]

Chrome 23.0.1271.95, on the other hand, modified timestamps for many files when terminating incognito sessions. As with Chrome 7.0.517.41, artifacts related to the user's browsing activities were limited to the computer's memory. However, the pagefile did reveal evidence of a file downloaded during an incognito browsing session.[4] The timestamp artifacts were still present as of version 25.0.1364.97.[6]

3.4 Microsoft Internet Explorer

Internet Explorer leaves the most artifacts behind when its InPrivate Browsing mode terminates.[4] Internet Explorer 8.0.6001.18702 continued to cache web pages on the hard drive during an InPrivate Browsing session, but these files were deleted once the browser window was closed. This means that fragments of these files—and thus a record of the user's browsing history—could be discovered by analyzing the free and slack space of the hard drive.[5] Login information for email accounts and full image artifacts were discovered on the machine after an InPrivate Browsing session using Internet Explorer 8.0 as recently as 2013. [4] Researchers recently determined that Internet Explorer log files no longer contained evidence of browsing history as of version 10.0.9200.16521. This had been an issue in previous versions of the browser.[6]

3.5 Proof of Concept: Parsing SQLite Databases and Browser Memory

As an example of the persistence of some of these private browsing shortcomings, two Python scripts were written to accompany this paper. The first parses the SQLite databases of Mozilla Firefox to look for private browsing artifacts. There are two built-in artifact searches. It will print any cookies created within the last hour; if private browsing is currently in use, some cookies may show up here before being removed upon browser termination. Additionally, URLs that have a nonzero access count but an unspecified last visit time will be printed. It cannot be easily determined when these websites were visited, but it is likely that it was during a private browsing session. Upon manual review of the URLs retrieved from this database, URLs visited more recently during a private browsing session appeared at the bottom of the list. The SQLite Manager extension for Firefox was used to discover these possible artifact sources before a programmatic solution was implemented.

The second script retrieves the contents of the memory segments associated with a running Firefox or Chrome process. String literals in the browser's memory are then logged to an output file, separated by browser. Additionally, potentially interesting memory contents (currently URLs, but could easily be modified to include other items of interest) are printed to the terminal. This was accomplished using the WinAppDbg module.

4 Conclusion

With all the ways in which private browsing modes can leak information, it may seem as though they are completely useless. This is not entirely true. While private browsing sessions won't help anyone evade law enforcement, a forensic investigation, or the United States government, they are still good enough to fool most computer users. Nosy roommates, family members, or patrons at the public library are unlikely to go to the lengths discussed in this paper to determine what someone else was doing on the internet.

That said, private browsing modes don't do much to keep a user's identity private to the outside world. Users seeking this kind of anonymity should consider using the Tor browser, which employs onion routing to connect them to their destination server. In this type of routing, a random path from the user to the destination is constructed, and content is encrypted at each point in the path. Furthermore, each point (known as a relay or node) only knows from which relay any given piece of network traffic immediately came and the next relay to which it is going. It does not know the original source of the network traffic or its final destination. This is what allows users to keep themselves and their activities anonymous using Tor in a way that the traditional web browsers cannot.[7] Of course, this type of browsing can be (ab)used for a whole host of illicit activities, but those are beyond the scope of this paper.

Until Mozilla, Google, and Microsoft reduce the artifacts leaker by their respective browsers' private modes, users who are concerned about "local attackers"—those of the nosy roommate or family member variety who have physical access to their machine—can do their part to keep their private information secure. Random access memory, where private browsing artifacts have been found, is cleared when a computer is powered down. Shutting down the computer after a private browsing session can reduce or eliminate these artifacts. Hard drive artifacts, however, are nearly impossible to eradicate. Even if the files whose timestamps were updated upon termination of a private browsing session could be deleted without affecting the browser, their fragments would still remain on the hard drive. The moral of the story is that private browsing is nothing more than a nice illusion useful for fooling the computer-illiterate. At the end of the day, private browsing really isn't that private at all.

References

[1] Aggarwal, Gaurav, Elie Bursztein, Collin Jackson, and Dan Boneh. "An Analysis of Private Browsing Modes in Modern Browsers." Proceedings of the 19th USENIX Security Symposium, Wardman Park Marriott Hotel, Washington, D.C. 11-13 Aug. 2013. Web. 8

Dec. 2013. < http://crypto.stanford.edu/~dabo/pubs/papers/privatebrowsing.pdf>

[2] "Chrome Browser." *Chrome*. Google, 18 Nov. 2013. Web. 10 Dec. 2013. <http://www.google.com/intl/en/chrome/browser/features.html#privacy>.

[3] "InPrivate Browsing." *Microsoft Windows*. Microsoft, 10 Dec. 2013. Web. 10 Dec. 2013. <http://windows.microsoft.com/en-us/internet-explorer/products/ie-9/features/in-private>.

[4] Ohana, Donny, and Narasimha Shashidhar. "Do Private and Portable Web Browsers Leave Incriminating Evidence? A Forensic Analysis of Residual Artifacts from Private and Portable Web Browsing Sessions." IEEE CS Security and Privacy Workshops (SPW), The Westin St. Francis, San Francisco, CA. 23-24 May 2013. Web. 9 Dec. 2013. <http://www.ieee-security.org/TC/SPW2013/papers/data/5017a135.pdf>

[5] Said, Huwida, Noora Al Mutawa, Ibtesam Al Awadhi, and Mario Guimaraes. "Forensic Analysis of Private Browsing Artifacts." 7th International Conference on Innovations in Information Technology, Abu Dhabi, United Arab Emirates. 25-27 Apr. 2011. Web. 10 Dec. 2013. <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?tp=&arnumber=5893816>

[6] Satvat, Kiavash, Matthew Forshaw, Feng Hao, and Ehsan Toreini. "On The Privacy of Private Browsing - A Forensic Approach." Proceedings of the 8th International Workshop on Data Privacy Management (DPM '13), Royal Holloway, University of London, Egham, UK. 12-13 Sept. 2013. Web. 9 Dec. 2013. <http://homepages.cs.ncl.ac.uk/feng.hao/files/DPM13.pdf>

[7] "Tor Project: Overview." *Tor*. Tor, 7 Dec. 2013. Web. 10 Dec. 2013. <http://www.torproject.org/about/overview.html.en>.

[8] Verdi, Michael et al. "Private Browsing." *Mozilla Support*. Mozilla Foundation, 29 Mar. 2013. Web. 10 Dec. 2013. <http://support.mozilla.org/en-US/kb/private-browsing-browse-web-without-saving-info>.