CCF Research Proposal
Enhancing ZFS Forensics Tools

J.L. van den Berg* and G.K. Kozar†
Department of Computer Science
University of Amsterdam, The Netherlands
*j.l.vanden.berg@vu.nl, †g.k.kozar@student.vu.nl

I. Introduction

In recent years, storage technology has rapidly evolved. Where in the 1980's we just broke the 1GB mark, todays commodity computers can easily accommodate terabytes of storage. While the storage media still changes rapidly, file systems have continued to adhere to the decisions made at the time of their creation.

Such reasons prompted Sun Microsystems in 2005 to develop a new kind of file system, named *ZFS* (originally: the Zettabyte File System), with the goal of providing strong data integrity, simple administration and immense capacity [1]. The unusual structure and operation of ZFS, mostly due to their excessive use of pointers and blocks, means that many existing forensics tools and techniques cannot be used to analyze store utilizing ZFS, but the need for such tools and techniques have not changed.

II. Our Idea

Previous research has already shown that extracting digital forensics artifacts from disks utilizing ZFS (or *zpools*s more generally) is possible, but no research is available on attempting forensics on *zpool*s that utilize more advanced features of ZFS, such as deduplication and compression.

It is, of course, reasonable to assume that such features can, and are, used in the real world. This necessitates forensics tools that are capable to handling such systems. We therefore propose to investigate the feasibility of conducting file recovery and forensic time-lining on destroyed *zpool*s, that employ these features.

III. Requirements

Since *ZFS* was built to be available for as many possible configurations as possible, *zpool*s can be made from anything ranging from entire hard disks to even single files. Therefore, the only requirements for our proposed research are access to a running Linux system, including the ZFSonLinux kernel modules, and some storage medium.

IV. Ethical Implications

The need to perform digital forensics on hard disks, and as such, on file systems, is clear and ethically unambiguous: a lot of work has already been done on this field. Our work does not come with any additional ethical implications.

V. Previous Research

There are a few papers published on the topic of ZFS forensics.

The papers *Digital forensic implications of ZFS* [2] and *Zettabyte File System Autopsy: Digital Crime Scene Investigation for Zettabyte File System* [3] are the two papers from 2009 pioneering research into this field. They cover how the nature of ZFS can be used to perform digital forensics, and as such will likely prove the back-bone of our research.

Forensic Timeline Analysis of ZFS [4] from 2014 is the most recent publication in this topic that we could find. This paper also provides an excellent summary of all literature available on ZFS forensics and data recovery.

References

[1] J. Bonwick, M. Ahrens, V. Henson, M. Maybee, and M. Shellenbaum, "The zettabyte file system," in *Proc. of the 2nd Usenix Conference on File and Storage Technologies*, 2003.

[2] N. L. Beebe, S. D. Stacy, and D. Stuckey, "Digital forensic implications of zfs," *digital investigation*, vol. 6, pp. S99–S107, 2009. [Online]. Available: \url{http://www.sciencedirect.com/science/article/pii/S1742287609000449}

[3] A. Li, "Zettabyte file system autopsy: Digital crime scene investigation for zettabyte file system," 2009. [Online]. Available: \url{http://clt.mq.edu.au/~rdale/teaching/itec810/2009H1/WorkshopPapers/Li_Andrew_FinalWorkshopPaper.pdf}

[4] D. Leigh and H. Shi, "Forensic timeline analysis of zfs." [Online]. Available: \url{http://research.dylanleigh.net/zfs-bsdcan-2014/zfs-timeline-paper.pdf}