

Socarrat: Building Cost-Effective Secure WORM Devices Following the Reverse File System Approach

Gorka Guardiola Múzquiz
Enrique Soriano-Salvador
Universidad Rey Juan Carlos, Spain

Introduction & Motivation. WORM (Write Once Read Many) devices allow data to be written only once, after which it can be read an unlimited number of times. These devices are used for logging and other purposes (like backups) and are essential for a wide range of applications. A recurring theme in data regulations is the requirement for regulatory-compliant storage that provides WORM guarantees, ensuring the data's retention, secure deletion, and compliant migration [9]. Additionally, data should be tamper-evident: any manipulation must be detectable by an auditor. Although it may appear straightforward, implementing secure WORM devices is highly complex. The traditional approaches use continuous feed printers [3] optical devices [7], content addressed storage [8] or specially designed hardware not available for the general public [5]. There are also numerous distributed approaches to address this problem (e.g. [6]).

In addition to these considerations, the possibility of a cyberattack on the machine storing the data must be addressed. If the attacker takes control of the system, she should not be able to delete or modify the WORM data of the log. At this point, solutions relying on file system capabilities are vulnerable: if the attacker elevates privileges, she can change the file system configuration, modify or delete the files, tamper with data blocks (even directly on the block device), change the addresses of content-addressed data blocks, or format the file system. Distributed systems are also vulnerable: a denial of service attack compromises the ability to write the to the logs at critical moments.

We propose Socarrat, a radical and cost-effective solution to address this problem locally with a simple external USB device: a small single board running Linux with USB OTG support. For example, a Raspberry Pi with Socarrat can export a 1 TB USB mass storage WORM device, which can be mounted on any regular operating system as an ext4 or exFAT file system (i.e. without running special software), for approximately \$100.

Socarrat is based on a novel if somewhat contorted approach: The *Reverse File System*. This approach involves inferring the file system operations occurring at higher layers by examining the blocks written to the storage device. Note that what is being exported is a block volume, which will be mounted in the user's machine as a regular file system. Socarrat focuses exclusively on write operations at the end of the logs, disregarding any other operation. Note that the logs are regular append-only files.

Example Scenario. Alice connects a USB black box (i.e. Socarrat running on a Raspberry Pi) to her server. The server's operating system detects a USB mass storage device formatted as an ext4 or exFAT file system, which is then mounted. Inside the mount point, there is a file named log (there can be more than one, we use this one as an example). The applications running on the server are able to use this volume as a regular one, performing the traditional system calls to work with the files. The only difference is that when the applications access the log file, only read operations and append-only write operations are effective; all other write operations on it are discarded. Later, Alice can extract the log file from the device, along with an additional file that authenticates all the entries added to the log during this period. Using a diagnostic tool, she or a third party (the auditor) can verify that the log data corresponds to the entries made during the operational period, ensuring that no records have been deleted or tampered with.

Research Problem. It is worth noting that building a secure WORM is not a trivial problem. Within the USB black box, the volume served to Alice's machine cannot be mounted locally to check for differences and discard invalid writes, as this would lead to significant concurrency problems: the volume would be mounted in two systems at the same time (i.e. in the Raspberry Pi and the user's machine).

Therefore, we must follow the reverse file system approach: we need to carefully analyze the block updates and infer the write operations being performed on the data structures and blocks of the file system.

This process is indeed complex. When a block write occurs in an ext4 file system, we need to determine whether it modifies the inode or the blocks (whether they are data blocks or intermediate data structure blocks like extents) of one of the log files for which we give WORM guarantees. If it does, we must check the integrity of the tree and make sure we are not in the midst of a partial update. In such case, we can infer the data of a write operation potentially coalescing multiple writes into one. While journaling can assist in this process, it also adds another layer of complexity.

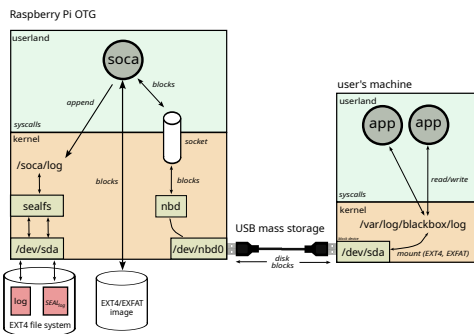
The Continuous Printer Model. The threat model is two-prone. First, a remote attacker, after a remote logical attack, can control the user's machine. Our goal in that case, is to minimize the attack surface to modify the data already stored in the log file: it is restricted the USB link. If the link

is not compromised, the logs are safe. Second, if the Socarrat system is compromised (by remote execution through the USB link or by physically accessing the external device), any tampering must be detected. To provide such forward-integrity properties, we use a previous system: SealFS [4, 10].

One of the challenges is precisely modeling the guarantees. Given the integrity of the USB link (which can only be used as a mass storage device), the guarantees provided by Socarrat are formalized by what we have named the The Continuous Printer Model (CPM):

1. Once committed, data cannot be deleted or rewritten.
2. Liveness: we commit data frequently enough while the system is operational, ensuring that writes do not wait indefinitely.
3. While spurious or bad data can be committed in the future if the client system is compromised, guarantee 1 is always preserved: what is *printed* cannot be *unprinted*. The same applies to other possible consequences of an attack (disk exhaustion, crashes, etc.).

Architecture. This is the general architecture of our system:



- Socarrat is the main component, written in Go, that implements the reverse file system approach.
- The nbd protocol is used to export the block device over the OTG USB link.
- A file system image (ext4 or exFAT) is used as backend for socarrat. It contains the file system exported to the user's machine.
- SealFS [4, 10] is the component that provides tamper-evident mechanisms for the protected log.

Related work. The idea of a reverse file system is quite unconventional and there is very little in terms of comparable systems or devices. The closest two we found are related to one another (one was inspired by the other). The first one is vvfat [2]. Vvfat is a block driver served by qemu which serves a virtual VFAT file system in which the files represent the files of an underlying directory and follow them when read and written. Inspired on this, there is the nbdkit floppy plugin [1]. These two devices implement a kind of *synthetic reverse file system*. The two main differences with Socarrat are:

1. These systems are unconstrained by a liveness property. They only need to guarantee that the underlying files are completely synchronized when they are unmounted. In contrast, our reverse file system must continuously extract valid operations to function properly.
2. The second is an implementation difference. A *reverse file system* observes the data structures present in a block device image. A *synthetic reverse file system* invents these data structures on the fly.

Performance. Note that our system is designed just for secure logging, it may not be suitable for high performance IO purposes.

To measure the if and how well the system works, we have the following challenges to meet:

1. The system is correct and works well on different operating systems when not under attack, and fails gracefully under attack (under the CPM constraints).
2. The system is sufficiently fast to be able to write logs at the speed needed for different applications.
3. We preserve the liveness guarantee and there is no bottleneck on the system which stops the logs from getting updated.

Currently, we have a working prototype on a Raspberry Pi 4 and are in the process of testing and measuring it.

References

- [1] Nbd virtual floppy disk from directory. URL <https://libguestfs.org/nbdkit-floppy-plugin.1.html>. [Online; accessed jan-2025].
- [2] Qemu block driver for virtual vfat. URL <https://github.com/qemu/qemu/blob/master/block/vfat.c>. [Online; accessed jan-2025].
- [3] M. Bellare and B. S. Yee. Forward integrity for secure audit logs. Technical report, University of California at San Diego, 1997.
- [4] G. Guardiola-Múzquiz and E. Soriano-Salvador. SealFSv2: combining storage-based and ratcheting for tamper-evident logging. *Int. J. Inf. Secur.*, 22(2):447–466, Dec. 2022. ISSN 1615-5262. doi: 10.1007/s10207-022-00643-1.
- [5] J. Leppäneniemi, T. Mattila, T. Kololuoma, M. Suhonen, and A. Alastalo. Roll-to-roll printed resistive worm memory on a flexible substrate. *Nanotechnology*, 23(30):305204, 2012.
- [6] M. Li, C. Lal, M. Conti, and D. Hu. Lechain: A blockchain-based lawful evidence management scheme for digital forensics. *Future Generation Computer Systems*, 115:406–420, 2021.
- [7] S. Quinlan. A cached worm file system. *Software: Practice and Experience*, 21(12):1289–1299, 1991. doi: <https://doi.org/10.1002/spe.4380211203>.
- [8] S. Quinlan, J. McKie, and R. Cox. Fossil, an archival file server. *World-Wide Web document*, 2003.
- [9] R. Sion. Strong worm. In *2008 The 28th International Conference on Distributed Computing Systems*, pages 69–76, 2008. doi: 10.1109/ICDCS.2008.20.
- [10] E. Soriano-Salvador and G. Guardiola-Múzquiz. SealFS: Storage-based tamper-evident logging. *Computers and Security*, 108:102325, 2021. ISSN 0167-4048. doi: 10.1016/j.cose.2021.102325.

Funding. This work is funded under the Proyectos de Generación de Conocimiento 2021 call of Ministry of Science and Innovation of Spain co-funded by the European Union, project PID2021-126592OB-C22 CASCAR/DMARCE.