

Fast and Secure Magnetic WORM Storage Systems

Yongge Wang and Yuliang Zheng
Department of Software and Information Systems
University of North Carolina at Charlotte
{yonwang, yzheng}@uncc.edu

September 7, 2004

Abstract

Computer forensic analysis, intrusion detection and disaster recovery are all dependent on the existence of trustworthy log files. Current storage systems for such log files are generally prone to modification attacks, especially by an intruder who wishes to wipe out the trail he leaves during a successful break-in. In light of recent advances in storage capacity and sharp drop in prices of storage devices, as well as the demand for trustworthy storage systems, it is timely to design and develop fast storage systems that practically have no limit in capacity and admit “secure append-only” operations (namely data can only be appended to a storage device; once appended it can no longer be modified, and can be read out by authorized users only.) This paper discusses an approach to building secure append-only storage systems. It proposes a possible secure append-only storage architecture that could be used to detect and prevent deletion or modification by inside/outside attackers. A specific implementation of the architecture based on block device drivers and magnetic storage firmwares is also presented.

1 Introduction

Today’s cyber-infrastructure has numerous vulnerabilities, and it is unlikely that this situation will change significantly in the short term. Thus it is very important to design detection and prevention mechanisms within the present infrastructure to enhance its security.

Security measures require the existence of trustworthy log data about a system to be examined. Two well-known examples are computer forensics and intrusion detection, both of which rely on inputs from log files about system and user activities. A major challenge is presented by sophisticated attackers who could delete all the system logs after a successful intrusion. An inside attacker may also easily delete all the trails that she/he has left. It is becoming increasingly evident that there is a need to develop storage systems that are able to provide unalterable, permanent records on computer system and user activities for the entire life span of a system. Specifically, the benefits of creating a complete, unalterable and permanent record of all activities on a computer system (infrastructure) include at least the following:

- Advanced intrusion detection systems can use these records to detect attacks as they happen.
- Computer forensic analysis could be carried out based on these records; the property of unalterability may also serve as court evidence in legal proceedings.
- The very presence of a secure and unalterable append-only log storage itself may act as quite an effective deterrence to potential inside and outside attackers. When a potential attacker knows that her/his

activities will be recorded there permanently and he/she has no way to delete them, she/he may hesitate to launch an attack in the first place.

- Legal requirements. SEC Rule 17a-4 [15] requires that companies under its jurisdiction maintain accessible, secure business records, and have the ability to produce records quickly that meet stipulated audit criteria—the penalties can be severe for non-compliance. The Sarbanes-Oxley Act of 2002 [24] takes this accountability to the individual level, placing personal requirements on executives to endorse their business’s financial statements, and to ensure that the organization has the ability to audit its business records, including electronic communications. Liability could potentially include financial penalties, exclusion from managing public companies, and even imprisonment. The Health Insurance Portability and Accounting Act (HIPPA) is also having a significant impact on append-only storage systems for medical firms.

A major goal of this paper is to design a secure append-only storage system to fight against attackers who may have root privileges to a computer system and who may try to delete log files containing information about the attacks. The attacks could be from either inside or outside attackers.

In what follows we discuss major requirements of storage systems for the creation of a permanent, immutable record on computer system activities, which is followed by the description of a secure append-only storage architecture that meets the requirements.

2 Requirements and related works

A storage system for recording computer system activities should satisfy the following requirements:

- Efficiency. It can record a huge amount of data on activities in real time and well before an attacker has time to delete them.
- Append-only. Even a sophisticated attacker should not be able to change any information that has already been written to a storage device.
- Permanency. The storage device should be able to provide permanent evidence for computer system and user activities.
- Unlimited capacity in practice. Each append-only storage device should be able to record activities for a relatively long period of time (e.g., a week), and an administrator can easily (and not so frequently) install a new storage device when the old one is full.
- The device should be cheap enough so that large scale deployment is economically viable. Furthermore, it is preferable that the secure append-only storage system could be built from COTS components.
- Security. The devices may contain critical and confidential information about computer activities. Thus the data stored on the device should be protected with strong cryptographic techniques in such a way that data can be readily accessible to authorized users only.

WORM (stands for write-once-read-many) storage devices could be used to fulfill parts of the above goals. Williams [25] classifies the WORM technologies into three categories:

- **P-WORMs** Physical WORM. For example, a CD-R is a P-WORM.

- **E-WORMs** Coded WORM. This includes driver level (embedded code) and media level (pre-recorded code) WORMS. For example, StorageTek's VolSafe™ technology is a driver level WORM.
- **S-WORMs** Software WORM.

According to [20], the existing WORM technologies have the following problems:

...most existing WORM implementations are based on older media technologies with limited capacity, slow data throughput, and substantial management overhead. Existing WORM storage options are optical WORM platters, which each holds about 30GB of data, and WORM tapes, with each cartridge able to store 50GB of data (best storage numbers based on latest technology). Traditional WORM storage vendors have gotten around capacity limitations of individual media by implementing expensive, complex media library and jukebox solutions that house numerous media. However, this solution for capacity limitations creates a problem of ever-increasing management overhead for the volumes of full media removed from the library or jukebox.

Among these WORM technologies, optical WORM platters (P-WORMs) are very slow in writing and WORM tapes are secure in several aspects. Typical WORM tapes include StorageTek's VolSafe™ and AIT worm tapes [2]. According to [2], if one inserts the AIT-WORM tape cartridge into a non-WORM drive, then the WORM functionality is achieved due to the following reasons:

- A non-WORM drive does not recognize WORM-application ID F0h of the RMIC (ROM area). Thus the tape is ejected as unknown format. For a normal tape, the application ID of the RMIC is 0.
- No writing will be allowed even if the RMIX data could not be retrieved because no recording-enable hole of the AIT-WORM tape cartridge. The tape will be read only.
- If application ID of the RMIC can not be read, the drive will read application ID of the system area. Then the tape cartridge will be recognized as the WORM. However, the tape will be read-only due to the lack of recording-enable hole.
- If application ID of the system area also can not be read, the drive will recognize the tape as a non-WORM tape. However, the tape will be read-only due to the lack of recording-enable hole.
- If one inserts a WORM tape into a legal non-WORM cartridge, then the cartridge will recognize the tape as a WORM device by application ID of the system area of the tape. The the tape will be read only.

These assumptions are acceptable in most environments. However, a malicious attacker may create a special tape cartridge by re-writing the firmware so that she/he can use it to erase data on WORM tapes. A malicious attacker may also dig holes on WORM tapes to let AIT-WORM tape cartridge to erase data on them.

Recently, several industry vendors designed WORM storage systems based on magnetic technologies (E-WORMs or S-WORMs). For example, Network Appliance [20] launched SnapLock™ WORM storage systems based on their existing two products:

1. ONTAP™, a microkernel operating systems developed by Network Appliance for their network attached storage (NAS) systems such as FAS960; and
2. NearStore, a nearline storage solution built on ATA-based magnetic disk drives.

The authors have not found any public descriptions about the technologies used in the SnapLockTM WORM storage systems. Since SnapLockTM allows users to partition disk space on NearStore file servers as WORM storage capacity, we think that the WORM functionality is implemented in the ONTAP operating system (that is, it is an S-WORM). This will be different from the method that we will introduce in this paper.

EMC [10] launched Content Addressed Storage (CAB) Centera recently. Centera is essentially a WORM storage system based on Centera's CentraStarTM software operating environments. Similarly, the authors have not found any public descriptions about the technologies used in the Centera WORM storage systems. Based on the similar observations that CentraStarTM software operating environment implements the WORM functionality, we think that Centera is an operating-system-software-based WORM technology (that is, it is an S-WORM device).

Debiez, Hughes, and Apvrille [9] have presented a technology to design WORM devices by establishing a WORM module external to the hardware storage device. This is similar to our device driver based WORM technologies that we will present in this paper. However, it is different to the firmware based WORM technologies that we will present in this paper.

Apvrille and Hughes [4] present a threat model for WORM technologies and discussed security issues related to P-WORM, E-WORM, and S-WORM devices. In particular, they introduce a time stamped virtual WORM system to protect data security. This is an important contribution since classical WORM technologies focus on securing mechanisms that write information onto the media, but not to data itself.

Applications based on WORM devices have been extensively studied. For example, the log file system in [12, 13] requires so-called *log devices*: non-volatile, block-oriented storage devices that support random access for reading, and append-only write access. Another application of WORM devices is the "Advanced Packet Vault (APV)" project at University of Michigan, which designed network traffic recording application based on CD-ROM technologies. Basically, their prototype writes captured network packets to long-term CD-ROM storage using encryption for later analysis. Though the APV provides a simple solution for intrusion detection purpose and meets some of the requirements that we have mentioned above, it is not a satisfactory solution. In particular, the APV has the following shortcomings:

- It is very slow to write due to the use of CD-ROM devices.
- The capacity of CD-ROM discs is limited. Even for a moderately large installation, one may need to replace discs in a very short time period.
- In the APV architecture, the system receives network (encrypted) packets and assembles them on magnetic disk for subsequent writing to CD-ROM. If the machine is under attack, then the attacker may have sufficient time to delete the assembled information on the magnetic disk before it is written to the CD-ROM. Though it is possible to install the APV system in a hardened OS that has no Internet services and an outside attacker may have no chance to get the root access for the APV system, an inside attacker who has local access to the APV system could still delete the data before they are written to the CD-ROM.

This analysis shows that while APV presents an interesting experiment, it is still open to find a better solution to the permanent record problem for network activities.

File systems for special purposes have been implemented in several operating systems. For example, the Linux Intrusion Detection System (LIDS) group has implemented the LIDS system [26] which could achieve the following goals:

- Read-Only Files/Directory. Read only files mean that they do not allow any user including root users to modify it.

- **Append-Only Files/Directory.** Append-only files mean that one can only append bytes to the end of the file, no users (including root users) can do any other operations on the files.
- **Exception-Files/Directory.** The files are not to be protected. In some cases, one wants to protect the whole directory but also want some specified files to be unprotected, so one can define the files as exception and the directory as read-only.
- **Protection-mounting/unmounting.** When some filesystems are mounted after system boot up, one can disable any user, including root, from unmounting the filesystems.

These implementations are done by adding special flags to the inode of protected files/directories and are based on Virtual Files Systems (VFS) or File Systems (FS). Since no protection is done at the device driver level, any privileged user (e.g., a root user) can access the hard disk directly via device drivers (e.g., `usedd` command) and delete these protected files/directories. Thus this solution does not satisfy all the requirements we discussed earlier.

3 Recommended solutions

In the following, we propose a secure append-only storage architecture using modified magnetic disks. Magnetic disks are re-writable, and generally do not satisfy the requirements of append-only property. In our solution, we propose to achieve WORM property for magnetic disks using one of the following mechanisms.

1. Modify block device drivers for magnetic disks and make these disks append-only. As soon as data is written on a specific block on the magnetic disk, no one can modify that block. Special blocks on the magnetic disks will be reserved to keep track of which blocks have been written. This solution only requires changes in the operating systems and could be easily implemented.
2. Modify the firmware inside the magnetic disk controllers and make these disks append-only. Similarly, certain area on the disks will be reserved to keep information on which blocks have been written.

With the second solution, one needs to modify the firmware in the disk controllers to achieve the goal, this could be impossible in several situations (e.g., manufacturers may be reluctant to provide source codes for the firmware, or a user cannot overwrite the firmware in the controller). However, if manufacturers could convert their magnetic disks into append-only storage systems by installing append-only firmware in the controllers of disks, we will be able to have an enhanced level of security.

Cryptographic techniques will be used to guarantee that appropriate security policies are automatically enforced in append-only storage systems. The architecture of our scheme is described in Figure 1.

4 WORM storage systems based on block device drivers

For the purpose of network monitoring, the BSD Packet Filter (BPF) [19] could be used (with appropriate modification) to capture desired packets and write them to a file in the main memory. To monitor host activities, log messages generated by the kernel could be directly written to a file in the main memory. When encryption is desired, these files will be encrypted before being written to the main memory. The file size should be chosen appropriately to ensure that

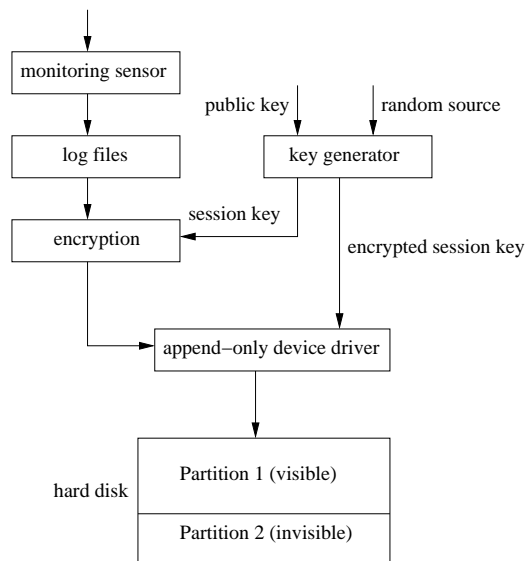


Figure 1: Architecture of append-only Storage Systems

1. Enough data could be collected in any specified time period for the file so that the file could be promptly written to the append-only storage system;
2. Files are large enough to avoid extensive disk fragmentation.

Conventional file systems are designed under the assumption that blocks on storage devices can be written more than once since file systems need to overwrite metadata frequently. This poses a challenge for the design of magnetic WORM storage systems that could be used in regular file systems. There are several ways to address this challenge.

- **Cached WORM file system [22]**. In this design, the file system resides on the WORM device. However, when a block of the WORM device is written, the data is not transferred directly to the WORM; instead, it is cached on an extra magnetic disk, the WORM cache, with the result that certain blocks can be written multiple times. The cached file system has been implemented in AT&T Bell Lab's experimental operating system Plan 9. Cached file systems could be used for our WORM storage devices. However, this may lead to compatibility problems since most operating systems on market do not support this file system.
- **File Motel [16]**. In this design, it uses a WORM to store backup copies of the files of conventional file systems. A separate database, which resides on an extra magnetic disk, is used to find the files on the WORM. The integrity of the WORM data is ensured by the ability to reconstruct this database if it is corrupted. This solution is feasible for our magnetic based WORM system only if it is possible to modify the file systems. Similarly, we have the compatibility problem.
- **Optical File Cabinet [14]**. In this design, it uses a block replacement strategy and place a general purpose file system directly on a WORM, taking into account the restriction that blocks can be written only once. The file system sees a logical address space that is much smaller than the size of the WORM device. Logical block addresses are mapped to physical WORM addresses; each time a logical block is written,

the mapping is changed so that an unwritten physical block is used. Thus a logical block can be written multiple times. The advantage of this solution is that file systems do not need to be changed and the functionality could be achieved by modifying device drivers.

- A further approach could be designed by modifying the File Motel method. That is, modify the magnetic device driver in such a way that the magnetic disk is divided into two parts. Part one of the disk serves as the extra magnetic disk in the File Motel method and part two of the disk serves as the WORM disk (the WORM functionality could be achieved in the way that we will describe).

In the following, we will present our WORM architecture and WORM implementation which is independent of file systems (e.g., NTFS, FAT, FAT32, ext2fs) that are used in the operating system level. The only changes to an existing system is done in the block device driver level. As we have mentioned above, in order for these magnetic WORM devices to be used in practice, special file systems need to be designed or special functionality need to be added to the device driver. Since our WORM solution is device driver (firmware) based and we prefer a file system independent solution, we assume that optical-file-cabinet-like solutions are used to address the metadata rewritability problem. That is, these functions are embedded in the device driver or firmware already.

4.1 Log file encryption and key management

Log files generally contain sensitive information about user activities, and should always be protected. After the log files are generated and written to the main memory, it should be encrypted before being sent to the append-only storage systems. Since these files are to be kept as permanent records, key management is a challenging issue. With different assumptions about the attackers, one could design different key management schemes. When designing such a scheme, one should be aware that if the attacker manages to get root privilege to the computer system, the attacker will be able to examine the content of the main memory and monitor the operation of the CPU. Thus the attacker may be able to get the encryption key.

Our scheme consists of public key cryptosystems and symmetric key cryptosystems. Public key cryptosystem is used for key management and symmetric key cryptosystem is used for log file encryption. The system administrators should first generate a public key and private key pair for a given public key cryptosystem. The private key should be kept in a secure place (threshold secret sharing schemes could be used to split the private key so that a certain number of persons is required to recover the private key). The computer system holds the public key. From time to time, the computer system generates a random session key for the given symmetric key cryptosystem, and encrypts the log files with this random session key. At the same time, the computer system encrypts the random session key, together with validity period and other auxiliary information, and writes this encrypted session key to the append-only storage system. In order to read the data on the append-only storage system, system administrators need to recover the private key for the public key cryptosystem first. The recovered private key can then be used to decrypt the random session key for the symmetric key cryptosystem. The data on the append-only storage system could be decrypted using this session key.

If the attacker gets root privilege at some time point, then the attacker may be able to find out the random session key used at that time point and the public key. From these information, the attacker could only read the log files encrypted with this session key (if the attacker can get these files). The attacker should not be able to recover the session keys that have been used to encrypt previous log files and the session keys that will be used to encrypt future log files.

If one has enough budget for tamper-resistant hardware, one can certainly install tamper-resistant hardware to hold the session keys and to encrypt log files. Thus the attacker will not even be able to get the session key

after getting root privileges.

4.2 Hard disk basics

A hard disk could be viewed as a continuous sequence of sectors, the smallest physical storage units on disks. For most disks, a sector is 512 bytes in size. Each disk sector has a factory tack-positioning label. Sector identification data is written to the area immediately before the contents of the sector and identifies the starting address of the sector.

When a disk is formatted, the most important information about the disk is written to the first sector (physical location: cylinder 0, side 0, and sector 1), which is normally called Master Boot Record (MBR).

The Master Boot Record contains the partition table for the disk and a small amount of executable code. For Intel x86-based computers, the executable code examines the partition table, and identifies the system partition (the location of the system boot sector). The MBR then loads the system boot sector into the memory and transfers the execution to the execution code in the system boot sector.

MBR consists of 446 bytes of the first sector of a disk (including the disk signature at the end of the MBR code) and 64 bytes of partition table. The partition table conforms to a standard layout that is independent of the operating system. Each partition table entry is 16 bytes long, making a maximum of four entries available. In another word, partition table 1 starts at the 446-th byte, partition table 2 at the 462-th byte, partition table 3 at the 478-th byte, and partition table 4 at the 494-th byte. The last two bytes of the first sector are a signature for the sector and are always 0x55AA.

An entry in the partition table consists of the following fields: boot indicator (one bytes, indicating whether the partition is the system partition. For example, on x86 based systems, 00 standards for non-bootable partition and 80 for system partition), starting head (one byte), starting sector (6 bits), starting cylinder (10 bits), system ID (one byte), ending head (one byte), ending sector (6 bits), ending cylinder (10 bits), relative sector (four bytes), and total sectors (four bytes).

4.3 Block devices and block device drivers

In most Operating Systems, hard disks are accessed as block devices via the block device drivers. A block device driver accesses the underlying device in multiple of the block size and usually allows random access. An important use of block devices is to support filesystems and swap files where the access is in multiple block units. I/O operations on files are first cached in the kernel's buffer cache before the device is invoked. Different operating systems may have different mechanisms for block devices. In the following, our discussion will be based on Linux operating systems.

Block device drivers are the media between operating systems and physical devices. Block device drivers present physical devices as continuous sequences of data bytes to the operating system.

A block device driver is a collection of routines that get called as various operations are performed on the devices controlled by the driver. The list of functions could include: `open()`, `release()`, `ioctl()`, `init()`, and `request()`.

I/O operations are always expensive. Thus block device drivers generally do not provide `read()` and `write()` functions directly. When a user process calls either of the `read()` or `write()` system calls related to the block device, the operating system will automatically handle it to the cache mechanisms.

When it is essential to write data to or read data from the block device, the buffer cache will add the I/O request to a queue of such requests for the corresponding device and then arrange for the `request()` function in the block device driver to be called to deal with the queue of requests.

The `request()` function reads each of the pending I/O requests in turn from the request queue and perform the physical read or write operations specified. For Linux systems, each I/O request is stored in a structure **struct request** defined in `/usr/include/linux/blkdev.h`. The general layout of a `request()` function in a device driver without an interrupt service routine looks like this:

```
static void do_request (void) {
    loop:
    INIT_REQUEST; /* make sure there
                  is at least one request */
    if (MINOR(CURRENT->dev) >
        MY_MINOR_MAX) {
        end_request(0);
        goto loop;
    }
    if (CURRENT->cmd == READ) {
        end_request(do_read());
        goto loop;
    }
    if (CURRENT->cmd == WRITE) {
        end_request(do_write());
        goto loop;
    }
    end_request(0);
    goto loop;
}
```

where `CURRENT` is a pointer to the **struct request** at the head of the request queue. The request structure layout is as follows:

```
struct request {
    int dev;
    /* physical device for this
       request */
    int cmd;
    /* command to perform (READ
       or WRITE) */
    int errors;
    unsigned long sector;
    /* sector number to start */
    unsigned long nr_sector;
    /* number of sectors to
       read or write */
    unsigned long current_nr_sector;
    char *buffer;
    /* kernel memory buffer for
       data read or written */
};
```

```

    struct semaphore *sem;
    struct buffer_head *bh;
    struct buffer_head *bhtail;
    struct request *next;
}

```

4.4 Achieving WORM storage systems via device drivers

Since all operating system accesses to hard disks are through the hard disk device drivers, it is possible to re-write the block device drivers so that files written to the disks cannot be deleted even by root users.

As we have mentioned in the previous section, a disk could be partitioned into a few partitions. In our scheme, we will use hard disks with two partitions (pre-partitioned on other computers). We will also assume that all bits on partition two is set to 0's at the beginning. Partition one will be used to store files and partition two will be used to record which sectors on partition one has been written before. In order to achieve the goal that any sector on partition one is written once, one can replace the READ and WRITE parts of the do_request function in the block device drivers with the following pseudo code.

```

    if (CURRENT->cmd == READ) {
        if (CURRENT->sector == 1) {
            ``read sector one into
            CURRENT->buffer and
            replace bytes between
            462 and 510 with 0x00``;
        } else {
            end_request(do_read());
        }
        goto loop;
    }
    if (CURRENT->cmd == WRITE) {
        ``check the bits between
        CURRENT->sector and
        (CURRENT->sector) +
        (CURRENT->nr_sector) on
        the partition two, if all
        these bits are 0's, then set
        first_time = 1, otherwise
        set first_time = 0``;
        if (first_time == 0) {
            goto loop;
        } else {
            ``write all 1's to the bits
            between CURRENT->sector
            and (CURRENT->sector) +
            (CURRENT->nr_sector) on
            the partition two``;
            end_request(do_write());
        }
    }

```

```

        goto loop;
    }
}

```

The pseudo code “read sector one into `CURRENT->buffer` and replace bytes between 462 and 510 with 0x00” in the READ parts of the code will filter out all information about partitions two, three, and four. In another word, any users (including root) will only see partition one on the disk.

In the WRITE parts of the code, the pseudo code “check the bits between `CURRENT->sector` and `(CURRENT->sector) + (CURRENT->nr_sector)` on the partition two, if all these bits are 0’s, then set `first_time = 1`, otherwise set `first_time = 0`” is used to determine whether the sectors requested for written have been written before. If some of these sectors have been written before (that is, some corresponding bits on the partition two are 1’s), then the written request is ignored. Otherwise, corresponding bits on the partition two are set to 1’s and execution is transferred to the code to finish the written request. In the above pseudo code, if any sector requested for written has been written already, the entire request will be dropped. One can also modify the pseudo code in such a way that the device driver will still write the sectors that have not been written before even if some sectors in the request have been written already.

In our above example, one bit on partition two is used to represent the status of one sector on partition one. Assuming that the sector size is 512 bytes (or 4096 bits), the capacity of partition two should be at least the $\frac{1}{4096}$ th capacity of the partition one.

The above discussion concentrates on the WORM implementation into the device drivers. In order for the WORM device to work in regular file systems, file-cabinet functionality [14] or modified file motel functionality needs to be embedded into the device driver also. Since the method used in [14] could be easily implemented within the device driver, we will not go to details.

4.5 Loop devices

A lookback device [5] in Linux is a virtual device that can be used like any other media device. A loopback filesystem associates a file on another filesystem as a complete device. The loopback devices can then be formatted and mounted as any normal block devices. To do this, the device called `/dev/loopi` is associated with the file and this new virtual devices is then mounted. The concept of loopback devices has been extensively used to achieve different goals. For example, it is used to encrypt block devices at sector level [23]. It has also been used in fibre channel Host Bus Adaptor (HBA) cards technologies. The advantage of loopback devices is that in order to achieve these goals, one does not need to hack the kernel codes. Loopback device technologies could also be used to produce append-only storage systems. However, this kind of append-only storage systems are not secure enough since any users with appropriate privilege could use low level commands such as `dd` to write the physical block device directly, thus bypassing the append-only protection mechanisms in the loopback device drivers.

Append-only storage systems based on the modification of block device drivers are more secure since even root users cannot write a sector two times.

4.6 Achieving WORM storage systems via firmware

In the previous sections, we discussed mechanisms to get WORM storage systems with device drivers. In order to achieve better security, one can combine these technologies with the controller firmware in the hard disks. For example, the firmware in the controller could make a reserved partition on the magnetic disk invisible to the users. This partition could be used to represent which sector in the visible partition has been written

already. The firmware will block any endeavor to write a sector if it has already been written already. In the meanwhile, file cabinet functionality should also be embedded into the firmware so that the WORM device could be used in regular file systems.

5 Comparison with existing magnetic WORM storage systems

As we have mentioned in the previous sections, currently there are mainly two magnetic WORM storage systems: Network Appliance's SnapLockTM and EMC's Centera. From the authors experience, we feel that these two technologies implement WORM functionality at the operating system level (that is, they are S-WORM devices). If this is a valid assumption, then our "block device drivers based WORM system" is approximately the same as these two technologies. The difference is that our system provides cryptographic protection of the data stored on the WORM storage systems. The virtual worm method and systems designed by Debiez, Hughes, and Apvrille [9] is essentially the same as our "block device drivers based WORM systems" though their system could be hardware or software based.

However, our "firmware based WORM storage systems" is quite different from these two existing technologies and from the virtual worm method and systems [9]. As we have mentioned in the previous sections, for block device driver based WORM storage systems, the disk volume itself has no protection when removed from host operating systems (we assume that EMC's Centera and NetApp's SnapLock have the similar disadvantages), a user could install the disk volume on another host machine and delete the data on it. For firmware based WORM magnetic storage systems, we have the further protection that one can not install the disk volume on another machine and delete the data contained in it.

In the previous sections, we have mentioned that a malicious attacker could attack the WORM tapes by creating his/her own tape cartridge or by digging holes on the WORM tapes. In certain sense, our firmware based WORM magnetic storage systems are secure against this kind of attacks since magnetic disks and firmwares are usually in one box and it is hard to produce a universal attacking box. Another obvious advantage of magnetic WORM storage systems is that one has random data access for reading and append-only writing operations.

6 Security considerations

In this paper, we introduced block device driver-based technologies to turn a normal magnetic disk into an append-only storage system. This system has the potential to be secure against not only outsider attacks but also sophisticated insider attacks. For example, an Internet attacker cannot delete the log files that he/she has left in the computer system since even if she/he has got the root privilege, she/he still cannot write a sector on the hard disk twice unless she/he could figure out the address in the main memory for the `do_request` function and modify the corresponding codes. It is generally extremely difficult to figure out these information for a specific running machine.

For dedicated attackers, they may monitor the operation of CPU and find out the location of codes for the corresponding device driver in the main memory. By modifying the codes for device drivers in the main memory, the attackers may be able to access the hard disk at a lower level and delete all data there.

An insider or an Internet attacker who has got the root privilege could certainly also install a new kernel onto the machine and reboot the machine. Then she/he will have a normal device driver and will be able to delete all data on the hard disk. We will assume that these kind of attacks will be hard to achieve. In cases that we have to address this kind of attacks, we recommend WORM magnetic storage systems based on firmware.

In the WORM storage system is based on block device driver technologies, an inside attacker could also walk into the server room, take the disk out, install it on another machine, and delete all data on it. An insider attacker could also copy the data on the WORM disk with modification to another WORM disk and destroy the old WORM disk. For these kinds of attacks, one may require that the disk be secured with deadlocks and at least two persons need to be present to take the disk out.

7 Implementation

A prototype of our magnetic WORM system has been implemented and tested by Feng [11] with Linux XFS file system. In the first implementation, only WORM functions are added to the device driver. When a magnetic disk was mounted with the modified device driver, it immediately crashes. the reason is that the file system need to rewrite metadata frequently. Then simple file cabinet [14] functionality was added to the device drivers. When a magnetic disk was mounted with this new device driver, it works perfectly as a WORM device. That is, even the root cannot delete data that has been written to the disk. We have not implemented cryptographic (key management) functionality yet, which should be relatively straightforward.

Acknowledgment

We would like to thank the referees for comments on this paper. In particular, we thank them for raising the issues regarding file systems over WORM storage systems.

References

- [1] L. Absher. VolSafe(TM): a discussion of non-erasable, non-rewritable tape for the business environment. White paper, Louisville CO, July, 2001.
- [2] Advanced Intelligent Tape. AIT-WORM overview, Tape Storage Solutions, Sony Electronics Inc. 2003. http://www.aittape.com/ait_worm.html.
- [3] C. Antonelli, M. Undy, and P. Honeyman. The packet vault: secure storage of network data. In: *Proceedings of the USENIX Workshop on Intrusion Detection and Network Monitoring*. April 9–12, 1999.
- [4] A. Apvrille and J. Hughes. A Time Stamped Virtual WORM System. In: *Workshop SECIO2 Securite de la Communication sur Intenet*, September, 2002, Tunis, Tunisia.
- [5] A. Bishop. The loopback root filesystem HOWTO. <http://www.linux.org/>.
- [6] R. Card, T. Ts'o, and S. Tweedie. Design and implementation of the second extended filesystem. In: *Proceedings of the First Dutch International Symposium on Linux* <http://e2fsprogs.sourceforge.net/ext2intro.html>
- [7] CITI. Projects: Advanced Packet Vault. <http://www.citi.umich.edu/projects/apv>.
- [8] Datalink white paper. An overview of WORM technology: sorting fixed-content non-erasable data, 2003. Available at: <http://www.datalink.com/whitepapers/WORM%20Technology.pdf>

- [9] J. Debiez, J. Hughes, and A. Aprille. Virtual WORM method and system. US Patent 6615330, September 2003.
- [10] EMC. EMC Centera Content Addressed Storage System, 2002. Available at: http://www.emc.com/pdf/products/centera/centera_ds.pdf.
- [11] Xiaogang Feng. Software magnetic WORM storage system with Linux XFS file system. Msc. Thesis. The university of Liverpool, August 2004.
- [12] R. Finlayson. *A Log File Service Exploiting Write-Once Storage*. PhD Thesis, Stanford University, July 1989.
- [13] R. Finlayson and D. Cheriton. Log files: an extended file service exploiting write-once storage. In: *Proc. ACM Symposium on Operating System Principles*, pages 139–148, ACM Press, 1987.
- [14] J. Gait. The optical file cabinet: a random access file systems for write once optical disk. *IEEE Computer*, May 1998.
- [15] General Rules and Regulations promulgated under the Securities Exchange Act of 1934.
- [16] A. Hume. The file Motel – an incremental backup system for Unix. In *Summer Usenix Conference Proceedings*, pages 61–7-, 1988.
- [17] T. Laskodi, B. Eifrig, and J. Gait. A UNIX file system for a write-once optical disk. In: *Summer Usenix Conference Proceedings*, pages 51–60, 1988.
- [18] G. Malan, D. Watson, F. Jahanian, and P. Howell. Transport and application protocol scrubbing. In: *Proceedings of the IEEE Infocom 2000 Conference*, Tel Aviv, Israel, 2000.
- [19] S. McCanne and V. Jacobson. The BSD packet filter: a new architecture for user-level packet capture. In: *Proc. of Winter USENIX Conf.*, pages 259–269, San Diego, 1993.
- [20] Network Appliance. WORM Storage on Magnetic Disks Using SnapLock™ on NearStore™ Appliances. 2003. Available at: http://www.netapp.com/tech_library/ftp/3263.pdf
- [21] T. Ptacek and T. Newsham. Insertion, deletion, and denial of service: eluding network intrusion detection. 1998.
- [22] S. Quinlan. A cached WORM file system. *Softw., Pract. Exper.* **21**(12): 1289-1299, 1991.
- [23] R. Rhea. Loopback encrypted filesystem HOWTO. <http://www.linux.org/>.
- [24] Sarbanes-Oxley Act. One Hundred Seventh Congress of the United States of America at the second session. 2002. Available at: <http://news.findlaw.com/hdocs/docs/gwbush/sarbanesoxley072302.pdf>.
- [25] R. Williams. P-WORM, E-WORM, S-WORM is a sausage a Wienie? Chicago, 1997.
- [26] H. Xie. LIDS Hacking HOWTO. Linux Intrusion Detection Systems, <http://www.lids.org/lids-howto/lids-hacking-howto.html>