

# Trail of Bytes: New Techniques for Supporting Data Provenance and Limiting Privacy Breaches

Srinivas Krishnan, *Member, IEEE*, Kevin Z. Snow, and Fabian Monrose

**Abstract**—Forensic analysis of computer systems requires that one first identify suspicious objects or events, and then examine them in enough detail to form a hypothesis as to their cause and effect. Sadly, while our ability to gather vast amounts of data has improved significantly over the past two decades, it is all too often the case that we lack detailed information just when we need it the most. In this paper, we attempt to improve on the state of the art by providing a forensic platform that transparently monitors and records data access events within a virtualized environment using only the abstractions exposed by the hypervisor. Our approach monitors accesses to objects on disk and follows the causal chain of these accesses across processes, even after the objects are copied into memory. Our forensic layer records these transactions in a tamper evident version-based audit log that allows for faithful, and efficient, reconstruction of the recorded events and the changes they induced. To demonstrate the utility of our approach, we provide an extensive empirical evaluation, including a real-world case study demonstrating how our platform can be used to reconstruct valuable information about the *what, when, and how*, after a compromised has been detected. We also extend our earlier work by providing a tracking mechanism that can monitor data exfiltration attempts across multiple disks and also block attempts to copy data over the network.

## I. INTRODUCTION

Today, postmortem intrusion analysis is an all too familiar problem. Our devices are repeatedly compromised while performing seemingly benign activities like browsing the Web [1], interacting on social-networking websites, or by malicious actors that use botnets as platforms for various nefarious activities [2]. Sometimes, threats can also arise from the inside (e.g., corporate espionage), and often lead to substantial financial losses. Underscoring each of these security breaches is the need to reconstruct past events to know *what* happened and to better understand *how* a particular compromise may have occurred. Sadly, although there has been significant improvements in computer systems over the last few decades, data forensics remains a very tedious process; partly because the detailed information we require to reliably reconstruct events is simply not there when we need it the most [3].

Loosely speaking, recent efforts in data forensic research have focused on tracking changes to file system objects by using monitoring code resident in kernel space, or by making changes to the application binary interface. However, without proper isolation these approaches are subject to tampering and therefore can not provide strong guarantees with respect to the integrity of recorded events. Malicious users can, for instance, inject code into either kernel or user space, thereby

undermining the integrity of logs maintained by the tracking mechanism. Virtualization [4] provides a potential avenue for enabling the prerequisite isolation criteria by providing a sandbox for operating system code and applications. For example, a hypervisor can mediate disk accesses at the block level by presenting a virtual disk to the virtual machine (VM). One disadvantage, however, is that this abstraction suffers from a “semantic gap” problem [5], in which the mapping between file-system objects and disk blocks are lost, thereby making it difficult to track objects beyond the disk layer.

In this paper, we propose an approach for monitoring accesses to data in a virtualized environment while bridging the semantic gap issue. Specifically, we provide an approach for monitoring accesses to data that originated from disk, and capture subsequent accesses to that data in memory—even across different processes. Our approach achieves this goal without any monitoring code resident in the virtual machine, and operates purely on the abstractions provided by the hypervisor. Operating at this layer mandates that we access the disk at the block layer, memory at the physical frame layer and system calls at the instruction layer—all of which offer substantial engineering challenges of their own. In that regard, our contributions are in the design and implementation of an accurate monitoring and reconstruction mechanism that collates and stores events collected at different levels of abstraction. We also provide a query interface for mining the captured information, and in doing so, provide forensic analysts with far more detailed information to aide in understanding what transpired after a compromise (be it a suspicious transfer of data or modification of files) has been detected. We also provide an extensive empirical analysis of our platform, including a real world case study.

The remainder of the paper is organized as follows. We first present some background and related work in Section II. Sections III and IV describe our design and architecture, including the various monitoring subsystems and the respective challenges with combining data from the various levels of abstraction. We discuss several threat models our framework helps protect against in Section V. Section VI describes how the framework may be used to protect data from inadvertent or intentional disclosures. In Section VII, we present a detailed empirical evaluation of the runtime overheads and accuracy of our logging and reconstruction techniques. To highlight the strength of our approach even further, we present a case study in Section VIII showing how the framework was used to uncover interesting forensic evidence from a laptop that had been connected to a public network for 2 weeks. We discuss attacks on, and limitations of, our current design in Section IX and conclude in Section X.

The authors are with the Department of Computer Science at the University of North Carolina at Chapel Hill, Chapel Hill, NC, 27599 USA. email: (krishnan,kzsnow,fabian)@cs.unc.edu

## II. BACKGROUND AND RELATED WORK

Generally speaking, computer forensics attempts to answer the question of who, what and how after a security breach has occurred [6]. The fidelity of the recorded information used in such analyses is highly dependent on how the data was collected in the first place. Keeping this in mind, the approaches explored in the literature to date can be broadly classified as either client-based approaches (that use application or kernel-based logging) or virtualization-based approaches (that use hypervisor based logging). While client-based approaches can provide semantic-rich information to a security analyst, their fidelity can be easily undermined as the logging framework is usually resident within the same system that its monitoring. Hypervisor-based approaches, on the other hand, are generally thought to lack the semantic detail of client-based approaches, but can achieve greater resistance to tampering, as the logging mechanisms reside in privileged sandboxes.

*a) Client-based Approaches:* File-system integrity and verification has a long history, with some early notable examples being the work of Spafford *et al.* on Tripwire [7] and Vincenzetti *et al.* on ATP [8]; both of which use integrity checks to verify system binaries (e.g., `/sbin/login`). Extending this idea further, Taser [9] detects unauthorized changes to the file-system and reverts to a known good state once malfeasance is detected. Solitude [10] extends this concept even further by using a copy-on-write solution to selectively rollback files, thereby limiting the amount of user data that would be lost by completely reverting to the last known good state. These systems do not record evidence on how an attack occurred or the data that was compromised. Instead they are geared primarily at efficient restoration back to a known good state.

More germane to our goals are systems such as PASS [11] and derivatives thereof (e.g., [12]) that provide data provenance by maintaining meta-data in the guest via modifications to the file-system. However, their approaches require extensive guest modifications and share the same problems of client-based systems.

*b) Virtualization-Based Approaches:* In order for virtualization-based approaches to work in a data forensic framework, they need to first overcome the disconnect in semantic views at different layers in an operating system [5, 13]. In particular, Chen *et al.* [5] provides excellent insight into advantages and disadvantages of implementing secure systems at the hypervisor layer. The challenges are generally related to performance and the difference in abstractions between the hypervisor layer and the guest virtual machine. While the issue of performance has been addressed as hypervisor technologies mature, the “semantic gap” still remains. Antfarm [14], Geiger [15] and VMWatcher [16] have bridged this gap for a given layer of abstraction, but no single work has tackled the problem of bridging the gap for a set of interconnected layers of abstraction (i.e., spanning disk, memory and processes) while preserving the causal chain of data movement.

Closely related in goals is the approach of King *et al.* [17] which provides an event reconstruction approach for relating processes and files. BackTracker reconstructs events over time

by using a modified Linux kernel to log system calls and relate those calls based on OS-level objects [18]. The semantic gap issue is bridged by parsing the memory contents of the virtual machine during introspection time using an event logger compiled with the virtual machine’s kernel headers. This approach is fragile, as any changes to the guest kernel will undermine their approach [18] [17]. Similarly, in their VM-based approach, it is not possible to monitor operating systems that are closed-source. While BackTracker made significant strides in this area, we find that relying solely on system calls to glean OS state has several drawbacks. For one, since it does not monitor memory events, data movements (such as a process sending a file over a network socket) can only be inferred as “potential” causal relationships; neither can it detect the exact object that was sent over the network. To be fair, these were not part of its stated goals. By contrast, the causal relationships we build attempts to capture access chains across processes, all-the-while storing the exact content that was accessed and/or modified.

Also relevant are the techniques used by Patagonix [19] and XenAccess [20] that employ forms of memory inspection. Patagonix’s goal is to detect changes between binaries on disk and their image in memory. XenAccess is positioned as an extensible platform for VM monitoring. Our goals and approach are different in that we use signals from different layers of the VM (i.e., the system-call, memory and storage layers) to correlate accesses to a monitored object. Lastly, this paper significantly extends our earlier work [21] to include new techniques for dynamic provisioning and selective blocking.

## III. DATA TRACKING

Our primary goal in this paper is to enable fast and efficient recording of events involving monitored data (e.g., a set of files on disk), at a granularity that allows a security analyst to quickly reconstruct detailed information about accesses to objects at that location. Conceptually, our approach is composed of two parts, namely an efficient monitoring and logging framework, and a rich query system for supporting operations on the recorded data. To support our goals, we monitor events to a collection of locations  $L$  (i.e., memory, disk or network) and record read or write operations on  $L$ . We denote these operations as  $O$ . Any additional operations (e.g., create or delete) can be modeled as a combination of these base operations. We tie these accesses to the corresponding causal entity that made them to ensure that a forensic analyst has meaningful semantic information for their exploration [22].

The approach we take to capture these causal relationships is based on an event-based model, where events are defined as accesses,  $O$ , on a location  $L$  caused by a some entity, i.e.,  $E_i(O, L) \rightarrow ID$ . Loosely speaking, an entity is modeled as the set of code pages resident in a process’ address space during an event. The distinct set of code pages belonging to that process is then mapped to a unique identifier. This event-based model also allows us to automatically record events that are causally related to each other, and to chain the sequences of events as  $\bigcup_i^n E_i$ . Intuitively, events are causally related based on the same data being accessed from multiple locations; i.e.,

we consider  $E_0(O, L)$  to be causally related to  $E_1(O', L')$  if the same data object resides in  $L$  and  $L'$ . The event model also facilitates the implementation of protection mechanisms based on realtime tracking of causally related events. We show one such example of data protection in this paper by blocking exfiltration attempts over the network.

Since the hypervisor views the internals of a VM as a black box, a key challenge is in realizing this model with minimal loss of semantic information. This challenge stems from the fact that the monitoring subsystem gets disjoint views of operational semantics at different levels of abstraction. For example, a read system call operates with parameters in virtual memory and the guest file system layer, which then spawns kernel threads that translate the file system parameters into blocks, after which the request is finally placed on the I/O queue. Without any code in the guest, the challenge is in translating these requests and chaining them together.

As we show later, one contribution of this work lies in our ability to link together the various events captured within the hypervisor. In what follows, we present our architecture and the design choices we made in building a platform that realizes the aforementioned model.

### A. Architecture

The monitoring framework is built on top of Xen [23] with hardware-virtualization [24]. At a high level, the Xen hypervisor is composed of a privileged domain and a virtual machine monitor (VMM). The privileged domain is used to provide device support to the unprivileged guests via emulated devices. The VMM, on the other hand, manages the physical CPU and memory while providing the guest with a virtualized view of system resources. This allows our framework to monitor—from the hypervisor—specific events that occur in the virtual machine. In order to support the framework we added  $\sim 5\text{KLOC}$  to the Xen hypervisor.

The framework is composed of three modules that monitor storage, memory, and system calls (see Figure 1). The modules are fully contained within the hypervisor with no code resident in the virtual machine. The system is initiated by monitoring accesses to a specific set of virtual machine disk blocks. The storage module monitors all direct accesses to these blocks and their corresponding objects, while subsequent accesses to these objects are tracked via the memory and system call modules. Specifically, the memory module in conjunction with the system call module allows the framework to monitor accesses to the object after it has been paged-in to memory, and also builds causal relationships between accesses. The memory module is also responsible for implementing the mapping function that lets us tie events to specific processes.

As a result of our design, each of these modules have to bridge the “semantic gap” prevalent at that layer of abstraction; i.e., blocks to files, machine physical addresses to guest virtual addresses, and instructions to system calls. Since the framework is built to log events happening in the guest, a single guest event might trigger multiple hypervisor events crossing various abstraction boundaries, e.g., consecutive writes to a file by a text editor will require disk objects to be mapped

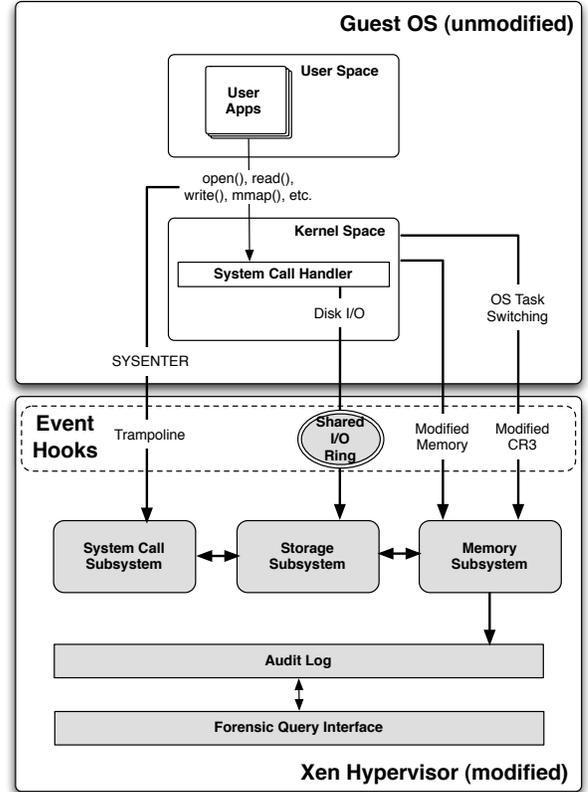


Fig. 1: Overall architecture of the forensic platform, depicting the memory, storage and system call layers

back to the file, writes to the page in the guest’s memory have to be mapped to the actual page in physical memory, etc. To effectively observe these linkages, our modules work in tandem using a novel set of heuristics to link events together. These events are stored in a version-based audit log, which contains timestamped sequences of reads and writes, along with the corresponding code pages that induced these changes. The log integrity is ensured using forward integrity hash chaining. We now turn our attention to the specific functionality of each of the monitoring modules.

1) *Storage Subsystem*: The storage module is the initialization point for the entire monitoring framework. That is, virtual machine disk blocks are monitored via a watchlist maintained by this module. Any accesses to the blocks on the watchlist triggers an update to the storage module. Accessing a block on the watchlist also notifies the memory module to monitor the physical page where the block is paged-in. In what follows, we first discuss how we monitor access at the block layer.

Figure 2 describes the Xen storage model and the enhancements we made to monitor disk I/O. In Xen, block devices are supported via the Virtual Block Device layer. Guests running on top of Xen see a virtual disk and therefore cannot directly modify physical disk blocks. Specifically, all accesses are mediated through the Xen storage layer, which exposes an emulated virtual disk. All I/O requests from the guest are written to an I/O ring, and are consumed by the storage layer.

The storage module monitors physical blocks on this virtual

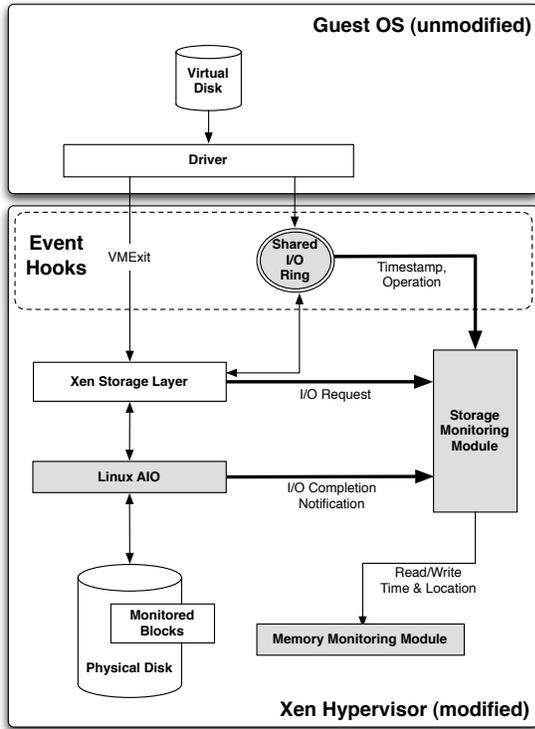


Fig. 2: The storage monitoring module, including our hooks for monitoring disk I/O at the Xen Storage and Linux AIO layers

disk and automatically adds them to watchlist it maintains. As guests place their I/O requests onto the shared ring, our monitoring code is notified via a callback mechanism of any accesses to blocks on the watchlist. This allows us to timestamp a request as soon as it hits the I/O ring—which is critical in matching the disk access with the system call that made the request, enabling the memory module to link a disk access with a specific process. Finally, the storage module waits for all reads/writes to be *completed* from disk before committing an entry in our logging data-structure.

a) *Dynamic Provisioning*: The aforementioned approach involving disk-level monitoring is not ideal, as tracking copies to unmonitored disks (the system disk, for example) would lead to tracking the entire destination disk as well. Since access to any block on a monitored disk triggers the storage module, each additional disk that requires monitoring could substantially increase the number of notifications that the storage module needs to process. Furthermore, it is usually the case that a given disk only has a limited set of files that require monitoring, whereas the disk level granularity does not allow us to exclude any given set of files on disk from being monitored. Therefore, in order to improve performance and precision, we extend the disk-level approach in [21] to explicitly disabled tracking accesses across unmonitored disks.

Instead, we introduce a new mechanism for tracking events system-wide, which we call *dynamic provisioning*. Rather than monitoring an entire virtual disk, we instead initialize the system with a set of file hashes that require monitoring. The blocks corresponding to the file hashes are added to the watchlist, and accesses to them trigger the storage module.

As derivations of the monitored data are created by user operations across different disks we *dynamically provision* the corresponding blocks onto our watchlist and monitor them as well. To do so efficiently, we leverage Xen’s storage model, which presents a consistent block-level device by abstracting away hardware level details. We also implement callbacks within Xen’s storage layer to receive notifications of data accesses across all the block devices available to the VM. The relevant blocks are then identified by querying the memory and system call module, a process that is described in the subsequent sections.

As we shall see later, the monitoring platform uses a special data structure to log all accesses. The amount of information we store in these logs is directly related to the number of blocks we monitor. By only tracking at the block level and not the entire disk, we are able to substantially reduce the number of blocks we monitor by allowing administrators to exclude files they deem non-essential, (e.g., static system files). Hence our technique of monitoring blocks on access and dynamically provisioning them has an added advantage of lowering the rate of growth of the logs.

As alluded by the discussion above, accesses to disk blocks typically happen as the result of a system call. In order to tie these two events together, it is imperative that we also monitor events at the system call layer. Next, we examine how we achieve this particular goal.

2) *System Call Monitoring Subsystem*: The system call module is responsible for determining when the guest makes system calls to locations of interest ( $L = \text{disk, memory or network}$ ), parsing the calls and inferring semantic linkage between related calls. We first describe how the module monitors the system calls and then discuss how they are used to infer semantic linkages with the memory monitoring module.

*Monitoring System Calls*: The use of hardware virtualization makes the efficient tracking of system calls in the guest an interesting challenge. To see why, notice that system calls on the x86 platform can be made by issuing either a soft interrupt  $0 \times 80$  or by using fast syscalls (i.e., `SYSENTER`). Modern operating systems use the latter as it is more efficient. This optimized case introduces an interesting challenge: a traditional  $0 \times 80$  would force a `VMEXIT` (thereby allowing one to trap the call), but fast syscalls on modern hardware virtualized platforms *do not induce a VMEXIT*. However, syscalls must still retrieve the target entry point (in the VM’s kernel) by examining a well-known machine specific register (MSR)<sup>1</sup>. Similar approaches for notification on system call events at the hypervisor layer have also been used recently in platforms like Ether [25].

Since the hypervisor sets up the MSR locations, it can monitor accesses to them. Our solution involves modifying the hypervisor to load a trampoline function (instead of the kernel target entry) on access to the MSR for syscalls. The trampoline consists of about 8 lines of assembly code that simply reads the value in `eax` and checks if we are interested in monitoring

<sup>1</sup>The `SYSENTER` call on the Intel platform uses the MSR `SYSENTER_EIP` to find the target instruction. This MSR is always located on Intel machines at address `176h`, even when setup by the hypervisor.

that particular system call before jumping into the kernel target point<sup>2</sup>. If we are, then the memory module (Section III-A3) is triggered to check the parameters of the call to see if they are accessing objects on the memory module’s watchlist. The trampoline code runs inline with virtual machine’s execution and does not require a trap to the hypervisor, avoiding the costly VMEXIT.

*Capturing the Semantic Linkage:* The system call module in conjunction with the memory module is responsible for inferring the semantic linkage between a set of related calls, for example, a `read()` call on a file whose blocks we monitor and a subsequent socket `open()`, `write()` of the bytes to a network socket. In this paper, we selectively monitor types of syscalls that could yield operations in our event model.

Specifically, we monitor syscalls that can be broadly classified as involving (1) *file system* objects, e.g., file open, read, write (2) *memory resident* objects, e.g., mmap operations (3) *shared memory* objects, e.g., ipc, pipes and (4) *network* objects, e.g., socket open and writes. As described earlier the system call module will monitor these calls and parse the parameters. The approach we take to create linkages between such calls is straightforward: we simply examine the source and destination parameters to infer data movement. In the aforementioned example, the system call monitor will be triggered on each of the file `read()`, network socket `open()` and `write()` calls. Since the source parameter of the `read()` references a monitored page, the memory module notifies the system call module of the offending access, and also adds the corresponding page of the destination parameter (e.g., the buffer) to its watchlist. When the memory module is later triggered because of the write on a network socket, that access will also be returned as an “offending” access since it references a page that is now on the memory module’s watchlist. As a result, the system call module will connect the two calls and infer the semantic linkage. Unlike other approaches that attempt to infer causal linkages based on data movements, our platform is able to accurately and definitively link events that are causally related. We now discuss the specifics of how the memory module decides if a particular event is accessing a monitored object.

3) *Memory Monitoring Subsystem:* The key function of this module is to track accesses to monitored objects once they are *resident* in memory. Recall that the initial access to  $L$  on disk causes the storage module to notify the memory module of potential data movement. This access causes a page fault, as the object has not yet been paged into memory. Since Xen manages the physical memory and hardware page tables, the fault is handled by the hypervisor. Our memory monitoring module is notified of this fault via the callback placed in Xen’s shadow page table mechanism, and updates its watchlist with the machine physical page of the newly paged-in monitored object. For brevity, we omit system level details and provide only the essential details. Before we proceed, we simply note that Xen provides the VM with a virtualized view of the physical memory by performing the actual translation from guest physical pages to actual machine physical pages [23].

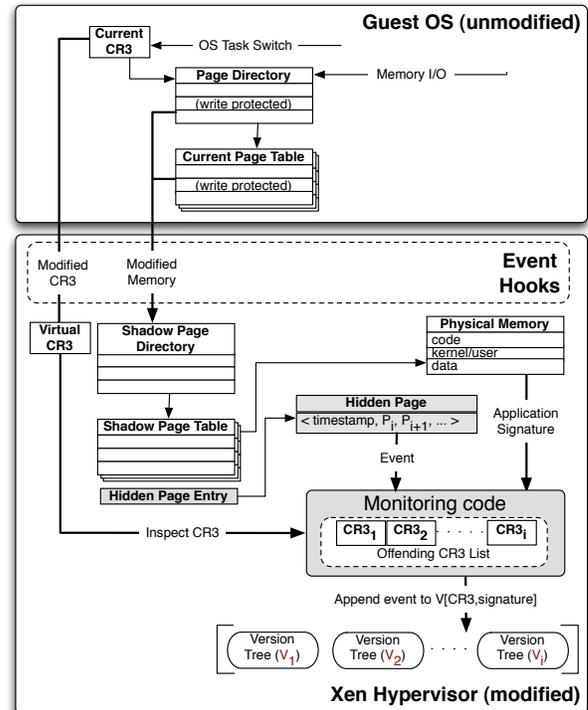


Fig. 3: Overview of the memory monitoring module, showing the hooks needed for tracking of monitored objects in memory and for logging the offending processes.

*Tracking objects:* The memory module uses its watchlist to track all subsequent accesses to monitored objects in memory. Recall that the system call module consults the memory module to determine if an access is to a protected object. To make this determination, the memory module consults its watchlist, and returns the result to the system call module<sup>3</sup>.

Notice that the memory monitoring module is in no way restricted to tracking only events triggered via system calls. Since it monitors objects in physical memory, *any* direct accesses to the object will be tracked. For instance, accesses to objects in the operating systems buffer cache will always trigger a check of the memory module’s watchlist.

Our approach extends the coverage of events even to accesses that might occur on monitored objects that are copied over to other memory locations. Since the memory monitoring module is triggered from the initial page-in event of the monitored data block from disk into memory, this paged-in machine physical page is automatically added to the watchlist. Hence, any subsequent events on this page such as a `mempcpy()` will result in the target memory location of the copy operation to be also added to the watchlist<sup>4</sup>. This is done to prevent evasion techniques that might copy the data into a buffer and then send the data over a network socket. Hence, any indirect data exfiltration attempts will also be recorded as an access to the original monitored block.

This is a key difference between the type of taint track-

<sup>3</sup>Recall the memory module must translate the guest virtual address to its physical address in a machine physical page.

<sup>4</sup>The destination machine physical page in `mempcpy`.

<sup>2</sup>System call numbers are pushed into `eax`.

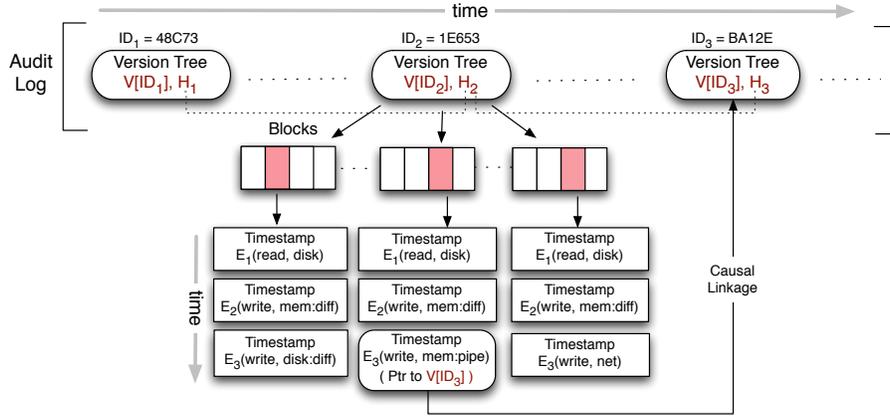


Fig. 4: The version tree stores different versions of blocks and the corresponding codepages that accessed these blocks over time. To support efficient processing of the audit log, we also store pointers to other version-trees of causally related processes.

ing [26, 27] commonly used to track objects in memory and the physical page monitoring we propose. Although taint tracking of that type affords for monitoring accesses to memory locations at a very fine granularity (e.g. pointer tracking), it does incur high overhead [28]. The memory tracking we implemented tracks accesses to the initial physical page frame where the data from monitored storage was paged in and subsequent physical memory locations the data was copied to. Our low overhead is achieved via a copy-on-write mechanism that tracks subsequent changes and accesses to the monitored objects. This implementation affords a coarser mechanism compared to taint tracking for memory monitoring, but achieves our goals at a much lower cost.

Once the decision is made that an access is to a monitored object, the memory module notes this event by timestamping the access<sup>5</sup>. The module also stores a “signature” of the code pages of the offending process. Recall that the CR3 register on the x86 platform points to the page directory of the *currently executing* process within the VM. Hence, to keep our overheads low, we do the signature creation lazily and add the address of the CR3 register (page-table register) to a queue of offending addresses that must be extracted later.

The signature is created as follows. For each item on this queue, we examine its page frames to inspect those codepages that are unique to the process being inspected. Because a CR3 could potentially point to different processes over time, we log the accesses in a modified B+-tree [29] where the root node is indexed by the tuple  $\langle \text{CR3}, \text{set of codepages} \rangle$ . In this way, we avert appending a new process’ events to an old process’ log. We call this structure a version-tree. The keys to the version-tree are the block numbers corresponding to the monitored object on disk, and the leaves are append-only entries of recorded operations on location  $L$ . The version-tree is built as follows:

- 1) If no version-tree exists for the process we are examining i.e. no tree has a root node that equals the current CR3 and code page hash, then let the set of known codepages be  $S = \emptyset$ , and skip to step (3).

<sup>5</sup>Specifically, a hidden page is appended in the shadow page table of the process with the timestamp and objects accessed.

- 2) Compare the hash of the codepages in the page table to the stored value in the tree. If the hashes are the same, there are no new codepages to record, and we only need to update the accesses made by this process; therefore, proceed to step (4).
- 3) To determine what new codepages have been loaded into memory, compute the cryptographic hash of the contents of the individual pages,  $c_i$ . Next, for each  $h(c_i) \notin S$ , determine whether it is a kernel or user page (e.g., based on the U/S bit), and label the page accordingly. If  $h(c_i)$  is found in page tables of more than one process, then label that page as shared.
- 4) Let  $S'$  be the set containing the hashes of user pages. Insert the access patterns (i.e.,  $E_0(O, L), \dots, E_1(O, L)$ ) into the version-tree with root node  $\langle \text{CR3}, S \rangle$ . That is, store the access time, location  $L$ , and “diffs” of the changed blocks for write operations, into the version-tree for that process. Update the root node to be the tuple  $\langle \text{CR3}, S \cup S' \rangle$ .

These version-trees are periodically written to disk and stored as an audit log where each record in the log is itself a version-tree (see Figure 4). Whenever the system call module notes a casual relationship between entities accessing the same monitored objects—e.g.  $E_i(O, L)$  by entity  $p_1$  and  $E_j(O', L')$  by  $p_2$ —we add a pointer in the version tree of  $p_1$  to  $p_2$ . These pointers help with efficient processing of the audit log.

Since logging is performed by a trusted entity in our framework – the hypervisor, our goal is to simply provide a mechanism that can detect data corruption. To that end the logging framework includes a check for testing individual entries for corruption. This is achieved by a hash-chain mechanism, whereby each log entry contains a hash element that verifies all other previous entries. Since every hash entry depends on previous entries detecting a corrupted entry is simply a matter of scanning all the hash chain entries within the log.

Having recorded the accesses to objects in  $L$ , we now discuss how the logs can be mined to reconstruct detailed information to aid in forensic discovery.

#### IV. MINING THE AUDIT LOG

To enable efficient processing of the data during forensic analysis, we support several built-in operations in our current prototype. These operators form our base operations, but can be combined to further explore the audit log. For the analyses we show later, the operations below were sufficient to recover detailed information after a system compromise.

- `report(w, B)`: searches all the version trees and returns a list of *IDs* and corresponding accesses to any block  $b \in \mathcal{B}$  during time window  $w$ .
- `report(w, ID)`: returns all blocks accessed by *ID* during time window  $w$ .
- `report(w, access, B | ID)`: returns all operations of type *access* on any block  $b \in \mathcal{B}$ , or by *ID*, during time window  $w$ .
- `report(w, causal, B | ID)`: returns a sequence of events that are causally related based on either access to blocks  $b \in \mathcal{B}$ , or by *ID*, during time window  $w$ .

##### A. Mapping blocks to files

Obviously, individual blocks by themselves do not provide much value unless they are grouped together based on a semantic view. The challenge of course is that since we monitor changes at the block layer, file-system level objects are not visible to us. Hence, we must recreate the relationships between blocks in lieu of file-level information. Fortunately, all hope is not lost as file-systems use various mechanisms to describe data layout on disk. This layout includes how files, directories and other system objects are mapped to blocks on disk. In addition, these structures are kept at set locations on disk and have a predefined binary format. As our main deployment scenario is the enterprise model, like Payne *et al.* [20] we assume that the file-system (e.g., ext3, ntfs, etc.) in use by the guest VM is known.

Armed with that knowledge, the storage module periodically scans the disk to find the inodes and superblocks<sup>6</sup> so that this meta-data can be used during forensic recovery. That is, for any set of blocks returned by a `report()` operator, we use the stored file-system metadata to map a cluster of blocks to files. For ease of use, we also provide a facility that allows an analyst to provide a list of hashes of files and their corresponding filenames. The `report()` operators use that information (if available) to compare the hashes in the list to those of the recreated files, and tags them with the appropriate filename.

#### V. THREAT MODEL

We have thus far described a forensic platform that transparently monitors and records data access events, follows the causal chain of these accesses, and supplies a tamper evident version-based audit log that allows for reconstruction of the recorded events. The platform may be used to thwart or analyze a couple of different threats, (a) *good-intentioned users*, and (b) *malicious programs*.

*Good-intentioned users* with full administrative privileges within the the guest OS running under a trusted hypervisor that may inadvertently leak sensitive information during their analyses of private data. The private data may, for example, be contained in a set of files in the user’s home directory. Sadly, such breaches of privacy occur all too often, as exemplified by the recent release of over 20,000 patient records from a Stanford Hospital. In that case, a spreadsheet with information including names, diagnosis codes, and billing charges, was attached to a question posted to a how-to forum, where the researcher was asking for assistance with data analytics<sup>7</sup>. Section VI discusses one way our introspection platform may be used to safeguard against this threat, while Section VII provides an empirical evaluation of how these safeguards may impact the user’s performance and the efficiency of our platform’s logging.

*Malicious programs* may also be inadvertently installed on a users system, typically through an exploit or trojan program. In this model, we assume the malicious program runs within the guest OS, while the hypervisor remains protected and trusted. These programs encompass a variety of behaviors, and may attempt to escalate privilege and install themselves directly into the running kernel (as with rootkits) or automatically collect and exfiltrate the user’s protected data. Section VIII exemplifies how the detailed audit logs generated by the introspection platform may be used for forensic analysis of this threat with several real-world examples. While the audit logs provide sufficient forensic information for many of today’s malicious programs, a clever adversary may be able to obscure their actions. Section IX discusses the potential attacks a clever adversary may use given knowledge of our platform and our overall limitations.

#### VI. DATA PROTECTION

While the discussion thus far has presented our introspection mechanisms as a technique for merely tracking and logging data accesses, the extensible design of our platform makes it relatively straightforward to extend its capabilities to thwart a good-intentioned user from inadvertently disclosing sensitive information. The desire to not only monitor, but to also block attempts to transmit data originating from a restricted datastore is very natural in security-sensitive scenarios.

To help minimize the risk of data exfiltration, we introduce a *selective blocking* mechanism. Selective blocking limits data exfiltration over the network by blocking the packets containing protected data and then notifying users of potentially accidental data exfiltration; it does so in real-time as the user performs the operation triggering the data exfiltration. We anticipate that the ability to selectively block connections containing sensitive data (i.e., in this case, data originating from a monitored store) would be of tremendous value to secure cloud computing environments for sensitive medical data. Indeed, our own university has engaged in an effort to provide medical researchers at the North Carolina Translational and Clinical Science Institute access to medical records hosted in private

<sup>6</sup>Similarly, the Master File Table and Master File Records under NTFS.

<sup>7</sup>See, for example, the NYTimes article entitled “Patient Data Posted Online in Major Breach of Privacy”, September 8, 2011.

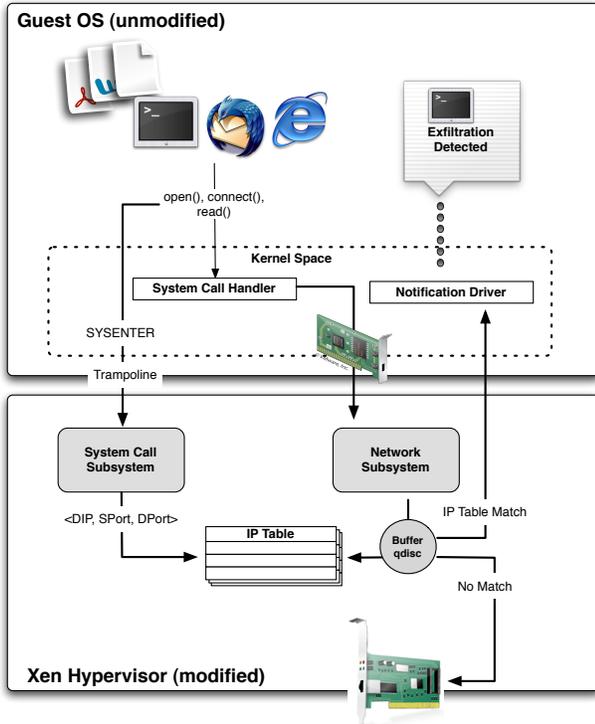


Fig. 5: Overview of the selective blocking mechanism, showing the blocking of packets using a custom queuing disc and the notification mechanism to the user

clouds, and it is this need that motivated the development of our selective blocking capability<sup>8</sup>.

Figure 5 describes the process to block packets that contain monitored data. Recall that the system call module traps specific system calls made by applications inside a virtual machine, parses the argument list and determines whether the call references a monitored memory location by consulting the watch list. We leverage this mechanism to identify potential exfiltration attempts by extracting the destination IP address and ports from network specific system calls (e.g. `connect()`, `sendto()` in Linux) if they reference a protected location in memory. The IP addresses and ports are then hashed and stored in an internal IP table within the hypervisor. The hash is used to identify network connections that reference data in protected memory and tag the packets by simply looking at the packet IP headers.

To block the packets, we implemented a network module that routes all packets from the virtual network interface of the VM to a custom queuing disc within the hypervisor. Linux supports implementation of custom network schedulers that can be registered for processing ingress and egress packets using what are known as *queuing disciplines* or *qdiscs*. The *qdisc* that we implemented inspects all egress packets and filters packets based on the rules in the internal IP table. The filtered packets are then stored within a buffer (per connection) within the hypervisor, the audit log is queried to find the offending process’s PID, and a notification is sent to the VM.

To interact with the user, an interrupt is injected into the

virtual machine by the network monitoring module when data is buffered. Within the virtual machine, a Windows notification driver handles the interrupt from the network monitoring module. This driver presents the user with a pop-up notification message containing information about the file name of the protected data being leaked along with the name of the offending application.

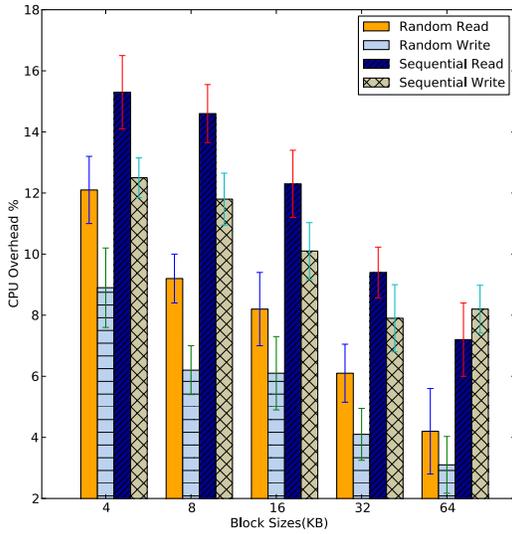
Our implementation of selective blocking incurs a fairly low overhead compared to data leak protection (DLP) systems that employ deep packet inspection (DPI) by performing fuzzy hashing on a packet’s payload. Unlike such heavy weight and error-prone approaches, packets are buffered by simply peeking at the IP header of each packet. Furthermore, DPI based solutions would be incapable of blocking connections that contain encrypted payloads (e.g. HTTPS connections), but we are not subject to this limitation as our solution works purely on the IP headers. We present an evaluation on this extension in Section VII-A

## VII. EMPIRICAL EVALUATION

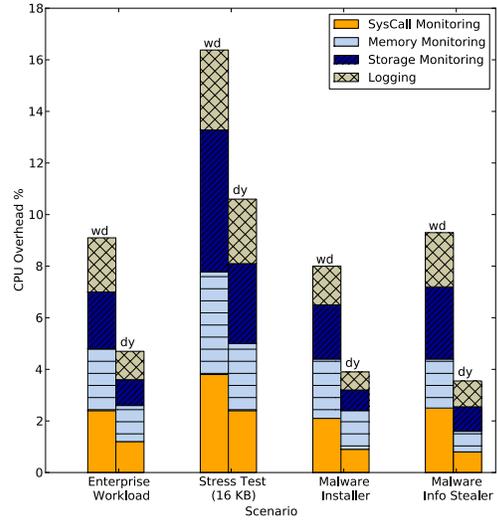
While having the ability to record fine-grained data accesses is a useful feature, any such system is impractical if the approach causes a high overhead. In what follows, we provide an analysis of our accuracy and overhead. Experiments were conducted on a 2.53GHz Intel Core2 Dual Core machine with 2GB of memory and Intel-VT hardware virtualization support enabled. Our modified version of Xen 3.4 with HVM support served as the hypervisor, and the guest virtual machines were either Windows XP (SP2) or Debian Linux (kernel 2.6.26). The virtual machines were allocated with 512MB memory, 1 virtual CPU, and with the hypervisor and virtual machine pinned to two different physical cores. We do so in order to reflect accurate measurements of CPU overheads. We mounted five disks – a 20GB system disk, two 20 GB data disks, a 10GB USB disk, and a 10GB network mapped disk. All disks contained both protected and extraneous files, and all experiments were run using *dynamic provisioning* mode (see Section III-A), unless otherwise stated.

First, the overhead associated with our approach was calculated under a stress test using a *Workload Generator* and a workload modeled for *Enterprise* users. Specifically, we subjected our design to a series of tests (using *IOMeter*) to study resource utilization under heavy usage, and used a scripting framework for Windows (called *AutoIt*) to automate concurrent use of a variety of applications. The application set we chose was Microsoft Office, plus several tools to create, delete, and modify files created by the Office applications. The parameters for the workload generator (e.g., the number of concurrent applications, average typing speed, frequency of micro-operations including spell-check in Word and cell calculations in Excel, etc.) were set based on empirical studies. The *Workload Generator* tests were conducted on an empty NTFS partition on one of the data disks, while the *Enterprise Workload* was tested with pre-seeded data comprising a set of Microsoft Office files along with additional binaries. These binaries performed various memory mapped, network and shared memory operations. The binaries were

<sup>8</sup>See *UNC’s Secure Cloud for Clinical Data* at <http://www.genomeweb.com>



(a) Varying block sizes and access



(b) Different Test Scenarios (wd = Whole Disk, dy = Dynamic Provisioning)

Fig. 6: Runtime overhead

added to increase the pool of applications loaded during the tests, and hence add greater diversity in the resulting code pages loaded into memory.

**Runtime Overhead:** Our runtime overhead is shown in Figure 6(a). The block sizes were chosen to reflect normal I/O request patterns, and for each block size, we performed random read, random write, sequential read and sequential write access patterns. The reported result is the average and variance of 10 runs. Each run was performed under a fresh boot of the guest VM to eliminate any disk cache effects. The IOMeter experiments were run on the same data disk with and without the monitoring code, and the overhead was calculated as the percent change in CPU utilization. The CPU utilization was monitored on both cores using performance counters. The reported utilization is the normalized sum of both cores.

Not surprisingly, writes have a lower overhead due to the increased time for completion from the underlying disk. Conversely, sequential access consumes more CPU as the disk subsystem responds faster in this case, and hence the I/O ring is quickly emptied by the hypervisor. Even under this *stress test*, the overhead is approximately 18%. This moderate overhead can be attributed to several factors in our design, including the scheduling of lazy writes of our data structures, the lightweight nature of our system-call monitoring, and the efficiency of the algorithms we use to extract the code pages.

Figure 6(b) shows a breakdown of CPU overhead across different test scenarios when using either dynamic provisioning or whole disk monitoring modes. As expected, dynamic provisioning significantly out-performs whole disk monitoring in all test scenarios. Notice that the majority of the overhead for the 16KB stress test scenario is attributed to the storage subsystem, as many of the accesses induced in this workload are for blocks that are only accessed once. We remind the reader that the expected use case for our platform is under the *Enterprise Workload* model, and the overall overhead

when using dynamic provisioning in this case is below 5%, with no single module incurring overhead above 1%. Also shown are the averaged overheads induced when monitoring and logging the activities of several real-world malware. In all cases, the overload is below 6%, which is arguably efficient-enough for real-world deployment. We return to a more detailed discussion of how we reconstructed the behavioral profiles of these malware using our forensic platform in Section VIII.

Another important dimension to consider is the growth of the log compared to the amount of actual data written by the guest VM. Recall that the audit log stores an initial copy of a block at the first time of access, and thenceforth only stores the changes to that block. Furthermore, at every snapshot, merging is performed and the data is stored on disk in an optimized binary format. We examined the log file growth by monitoring the audit log size at every purge of the version-trees to disk (10 mins in our current implementation). In the case of the *Enterprise Workload*, the experiment lasted for 1 hour, with a minimum of 4 applications running at any point in time. During the experiment, control scripts cause the overall volume of files to increase at a rate of at least 10%. The file sizes of the new files were chosen from a zipf distribution, allowing for a mix of small and large files [30]. We also included operations such as `make` to emulate creation and deletion of files. The overhead (i.e., additional disk space used to store logs and metadata compared to the monitored disk blocks) was on average  $\approx 2\%$ . Since the *Enterprise Workload* is meant to reflect day-to-day usage patterns, the low overhead indicates that this platform is practical and deployable.

**Accuracy of Reconstruction:** To examine the accuracy of our logging infrastructure, we explore our ability to detect accesses to the monitored data store by “unauthorized” applications. Again, the *Enterprise Workload* was used for

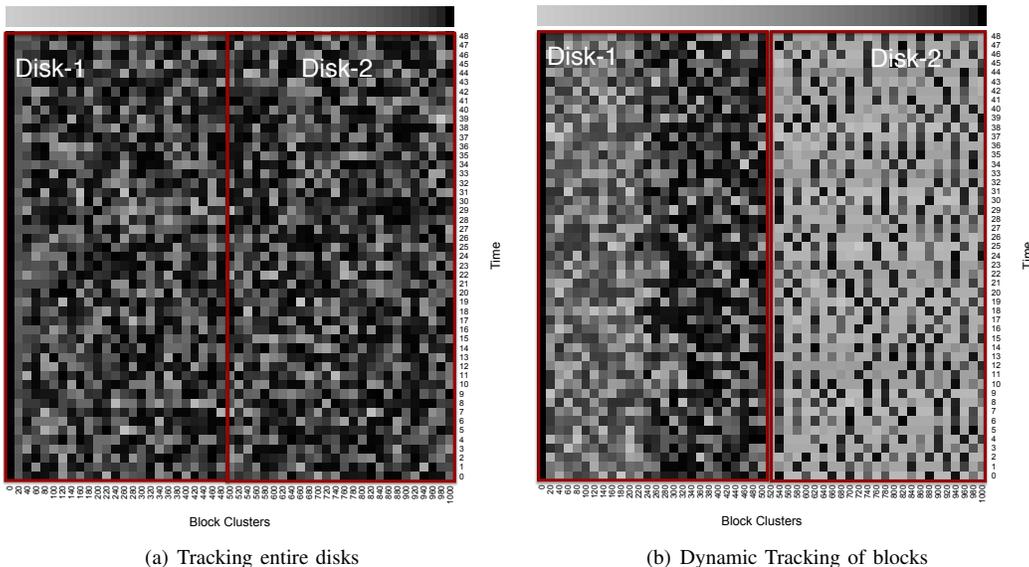


Fig. 7: Heatmap of block tracked over two days, darker shades indicate blocks that are tracked for multiple hours

Conn/min	VMEXITS ( $\mu s$ )	Insertion ( $\mu s$ )	Blocked %
10	1.2	0.5	100%
100	1.8	1.1	100%
1000	2.5	2.0	100%
10000	3.8	6.5	98.9%

TABLE I: Breakdown of time taken to insert a 5-tuple into the IP table and the accuracy of our blocking mechanism

these experiments, but with a varying concurrency parameter. Specifically, each run now includes a set of authorized applications and a varying percentage of other applications that also perform I/O operations on monitored blocks. The ratio of unauthorized applications for a given run was increased in steps of 5%, until all applications running were unauthorized. We also selected a set of files and provide their hashes to the *selective blocking* module to evaluate the effectiveness of the platform in preventing exfiltration of data. The task at hand is to reconstruct all illicit accesses to the disk. These illicit accesses include copying a file into memory, copying a file to the USB disk, sending a file over a network connection, and shared memory or IPC operations on monitored objects. To reconstruct the illicit accesses, the audit log was queried for the time-window spanning the entire duration of the experiment to identify both the unauthorized applications and the illicit access to blocks. We achieved a true positive rate of 95% for identification of the illicit applications and a 97% true positive rate in identifying the blocks accessed by these applications.

#### A. Dynamic Provisioning

To evaluate the performance benefits of tracking targeted blocks over an entire disk, we set up multiple disks on our network honeypot and tracked the movement of protected files over these disks. The set of files we placed for this experiment were chosen specifically to *bait* malware (e.g. W32.Tapin and W32.Pilleuz) that copy themselves over network and external drives and other infostealers.

Figure 7 shows blocks being tracked over the two data disks over two days, the darker shades indicate blocks that are tracked over multiple hours, whereas the lighter colored regions are blocks not being tracked. Under whole disk monitoring mode, as malware copies itself over to the secondary disk or attempts to copy files, all the blocks contained on the disk are automatically considered to be protected and hence tracked. This is evident from Figure 7(a) where we see dark regions on both disks, even though the secondary disk only contains a few protected files. However, when we employ *dynamic provisioning*, only blocks corresponding to protected files are constantly monitored. Figure 7(b) illustrates this with the presence of only a few dark (tracked) regions on the right side of the heatmap. Since we are only tracking specific disk regions, we see a major improvement in CPU performance as the number of VMEXITS decrease significantly under dynamic provisioning mode.

*Selective Blocking:* Selective Blocking requires a fast insertion of the network 5-tuple into the internal IP table in order to effectively block exfiltration attempts. To test the performance of the module under varying network load, we wrote a simple file transfer utility that creates a specified number of network connections and attempts to transfer protected files to an external machine. Iperf [31] was used to generate background traffic during the experiments.

Table I shows the breakdown of the average time taken to insert a 5-tuple into the internal IP table we maintain. During the experiment we gradually increase the number of unique connections that the utility generates from 10 to 10,000. As noted earlier, every insertion requires parsing of the arguments for specific systems calls and then extracting the 5-tuple. The information is then passed to the *selective blocking* module (see Section VI) via a VMEXIT. As we can see from the table, VMEXIT times are less than 4 microseconds, even in the worst case. The insertion into the internal IP table involves acquiring locks, hence the spike in insertion time

as we increase the number of connections. Since selective blocking mode prevents file transfers, the client application should ideally receive no packets. This is reflected in our blocking accuracy, where all packets are blocked except for a few in the 10,000 connection test. In this case, we were unable to block a few TCP SYNs due to queuing of insertion operations at the network module. However, we remind the reader that 10,000 unique connections/minute is extremely high.

### VIII. REAL-WORLD CASE STUDY

To further showcase the benefits of our platform, we report on our experience with deploying our framework in an open-access environment that arguably reflects the common case of corporate laptops being used in public WiFi environments. Specifically, we deployed our approach on a laptop supporting hardware virtualization, on top of which we ran a Windows XP guest with unfettered access to the network. The enterprise workload was configured to run on the guest system to simulate a corporate user. Similar to the earlier experiment, we had 5 disks attached to the machine – a system disk, two data disks, a network share and a USB disk. All the drives were seeded with files described earlier in Section VII-A and we also chose a subset of these files to provide to the *selective blocking* module. While there was no host or network-level intrusion prevention system in place on the guest system, we also deployed Snort and captured network traffic on a separate machine. This allowed us to later confirm findings derived from our audit mechanism. The laptop was left connected to the network for one week, and its outbound traffic was rate-limited in an attempt to minimize the risk of infecting other network citizens.

Syscall %	Phalanx2		Mebroot	
	Manual	Forensic	Manual	Forensic
Storage	72%	68%	91%	95%
Memory	26%	30%	8%	5%
Other	2%	2%	1%	0%

TABLE III: Comparison of the profiles created by manual analysis vs reconstruction using our platform

To automate the forensic recovery process, we make use of a proof-of-concept tool that mines the audit logs looking for suspicious activity. Similar to Patagonix [19], we assume the existence of a trusted external database,  $\mathcal{D}$ , (e.g., [32]) that contains cryptographic hashes of applications the system administrator trusts. The code pages for these authorized applications were created using a userland application that runs inside a pristine VM and executes an automated script to launch applications. The userland application communicates with the memory monitoring module, and tags the pages collected for the current application. The pages are extracted as described in Section III-A3, and are stored along with the application tags. Notice that these mappings only need be created once by the system administrator.

We then mined the log for each day using `report(24hr,  $\mathcal{B}$ )` to build a set of identifiers ( $p \in P$ ), where  $\mathcal{B} = \{\text{blocks for the temp, system, system32 directories and the master boot$

record}. Next, we extracted all causally related activity for each  $p \notin \mathcal{D}$ , by issuing `report(24hr, causal,  $p$ )`. The result is the stored blocks that relate to this activity. These blocks are automatically reassembled by mapping blocks to files using the filesystem metadata saved by the storage module (as discussed in Section IV-A). At this point we have a set of unsanctioned applications and what blocks they touched on disk. For each returned event sequence, we then classified it as either (i) an *info stealer*: that is, a process that copied monitored objects onto an external location (e.g.,  $L = \text{network}$ ) or (ii) an *installer*: a process that installs blocks belonging to an info stealer.

To do so, our recovery utility first iterates through the set of unsanctioned applications and checks the corresponding version-trees for events that match an info stealer’s signature. For each match, we extract all its blocks, and issue `report(24hr,  $b_i, \dots, b_n$ )`. This yields the list of all unsanctioned applications that touched an info stealer’s blocks. From this list, we searched for the one that initially wrote the blocks onto disk by issuing `report(24hr, write,  $b_i, \dots, b_n$ )`. The result is an installer.

Table II shows the result of running our proof-of-concept forensic tool on the audit logs collected from the laptop. The table shows the percentage of activity for each malicious binary and the classification as per the tool. For independent analysis, we uploaded the reconstructed files to Microsoft’s Malware Center; indeed all the samples were returned as positive confirmation as malware. We also subjected the entire disk to a suite of AV software, and no binaries were flagged beyond those that we already detected by our tool.

To get a better sense of what a recovered binary did, we classify its behavior as active if it had activity in the audit logs every day after it was first installed; or passive otherwise. The label “Exfiltration” means that data was attempted to be shipped off the disk. “Disk search” means that the malware scanned for files on the monitored store. As the table shows, approximately 75% of the recorded activity can be attributed to the info stealers. Upon closer examination of the blocks that were accessed by these binaries, we were able to classify the files as Internet Explorer password caches and Microsoft Protected Storage files. An interesting case worth pointing out here is Zeus. The causal event linkage by the forensic tool allowed us to track the initialization of Zeus as Zbot by Sinowal. Even though Sinowal constitutes only 4% of activity in the logs, it was responsible for downloading 60% of the malware on the system. Zeus appears to be a variant that used Amazon’s EC-2 machines as control centers<sup>9</sup>. Finally even though malware such as Kenzero, Tapin, Pilleuz were found to have low activity in the logs, they nevertheless had interesting behavior worth noting. Pilleuz and Tapin scanned the system for the presence of network and removable drives and would copy themselves over to those drives as a possible way to increase the spread of the worms. Pilleuz would also scan removable drives (e.g., our mounted USB drive), and scan for possible files to exfiltrate. Mebormi is a fairly recent malware that attempts to infect the Master Boot Record (MBR), similar to Mebroot that we perform detailed

<sup>9</sup>We verified this hypothesis independently based on our network logs.

Malware	% Activity in Log	Disk search	Exfiltration	Classification
Zeus & Variants	30.0	active	active	info stealer
Ldpinch	20.5	active	active	info stealer
Alureon	15.0	active	active	info stealer
Koobface	10.0	passive	active	installer
Bubnix	5.0	passive	active	installer
Masavebe	5.0	passive	active	both
Sinowal	3.5	active	active	both
Pilleuz	4.5	active	active	both
Tapin	3.0	passive	passive	installer
Mebromi	2.5	passive	passive	installer
Kenzero	1.5	active	active	info stealer

TABLE II: Malicious applications recovered from the audit log, and their high-level classification.

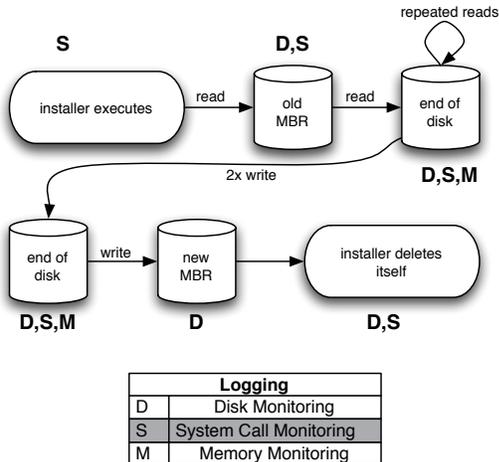


Fig. 8: Annotated graph of the causal reconstruction of Mebroot’s attack vector as recovered from processing the audit logs

analysis on in §VIII-A. Interestingly, the average growth of our audit log was only 9 MB per day compared to over 400 MB per day from the combined Snort and network data recorded during the experiment. Yet, as we show later, the data we are able to capture is detailed enough to allow one to perform interesting behavioral analyses. The analysis in Table II took less than 4 hours in total to generate the report, and our proof-of-concept prototype can be significantly optimized. Finally the *selective blocking* mode blocked all of the exfiltration attempts, as confirmed by the absence of alerts in the Snort logs.

#### A. Example Reconstruction

With the framework at our disposal, we decided to explore its flexibility in helping with behavioral analysis. Specifically, we were interested in analyzing Mebroot, which is a part of the stealthy Sinowal family. Mebroot serves as a good example as reports by F-Secure [33] label it as one of the “stealthiest” malware they have encountered because it eschews traditional Windows system call hooking, thereby making its execution very hard to detect. The anatomy of the Mebroot attack can be summarized as follows: first, a binary is downloaded and executed. Next, the payload (i.e., from the binary) is installed, and the master boot record (MBR) is modified. Lastly, the installer deletes itself.

To understand what Mebroot did, we issued `report( $\infty$ , causal, ID(Mebroot))`. The reason why the causal relationship

between the first two steps is built by our monitoring infrastructure should be obvious. In our platform, the connection between the first and last steps is made when the file deletion is noted (i.e., when the storage module rescans the inodes). An annotated profile of the behavior recovered from our audit log is shown in Figure 8. Notice that because we store “diffs” in the version trees, we are also able to see all the modifications made to the master boot record.

To further evaluate the strength of our platform in helping an analyst quickly reconstruct what happened after a compromise is detected, we provided two malware samples to a seasoned malware analyst (i.e., the second author) for inspection. In both cases, the malware was successfully unpacked and disassembled using commercial software and inspected using dynamic analysis techniques for system-call sequence analysis, for finding the payload in memory, and for single-stepping its execution. We then compared our results to those from this labor-intensive exercise.

The breakdown in terms of diagnosed functionality is shown in Table III. The overall results were strikingly similar, though the analyst was able to discover several hooks coded in Phalanx2 (a sophisticated info stealer) for hiding itself, the presence of a backdoor, and different modes for injection that are not observable by our platform. From a functional point of view, the results for Mebroot were equivalent. More important, however, is the fact that the manual inspection verified the behavioral profile that we reported, attesting to the accuracy of the linkages we inferred *automatically*.

## IX. ATTACKS AND LIMITATIONS

As stated earlier, the approach we take relies on the security properties of the hypervisor to properly isolate our monitoring code from tampering by malicious entities residing in the guest OSes. This assumption is not unique to our solution, and to date, there has been no concrete demonstration that suggests otherwise. However, if the security of the hypervisor is undermined, so too is the integrity and correctness of the transactions we record. Likewise, our approach suffers from the same limitations that all other approaches that have extended Xen (e.g., [16, 17, 20, 25]) suffer from—namely, that it extends the trusted code base.

A known weakness of current hypervisor designs is their vulnerability to hypervisor-detection attacks [34–36]. One way to address these attacks might be to rely on a thin hypervisor layer built specifically for data forensics, instead of using

a hypervisor like Xen which provides such a rich set of functionality (which inevitably lends itself to being easily detected). Once the presence of a hypervisor has been detected, the attacker can, for instance, change the guest VM's state in a way that would cause the forensic platform to capture a morphed view of the VM [36]. An example of such an attack would involve the attacker attempting to circumvent our event model by modifying the System Call Tables in Linux or the SSDT in Windows to remap system calls. This could cause false events at the system call layer and pollute the audit logs. That said, such an attack poses a challenge for all the hypervisor-based monitoring platforms we are aware of. Techniques to mitigate such attacks remain an open problem.

The techniques described in this paper depend on dynamic flow tracking of memory location. A known drawback of this approach are sophisticated attacks using implicit flows [37]. For example an attacker might perform a bitwise copy of the protected data using local variables. Currently our system does not prevent such attacks since prevention of such attacks require detailed static analysis.

Resource exhaustion attacks offer another avenue for hindering our ability to track causal chains. As our infrastructure tracks all monitored objects in memory, an attacker could attempt to access hundreds of files within a short period of time, causing the memory monitoring module to allocate space for each object in its watchlist. If done using multiple processes, the attack would likely lead to memory exhaustion, in which case some monitored objects would need to be evicted from the watchlist. While we have built several optimizations to mitigate such threats (e.g., by collapsing contiguous pages to be tracked as a single address range), this attack strategy remains viable. Lastly, since we do not monitor interactions that *directly* manipulate the receive and transmit rings of virtual network interfaces, such accesses will not be logged.

## X. CONCLUSION

We present an architecture for efficiently and transparently recording the accesses to monitored objects. Our techniques take advantage of characteristics of platforms supporting hardware virtualization, and show how lightweight mechanisms can be built to monitor the causal data flow of objects in a virtual machine—using only the abstractions exposed by the hypervisor. The heuristics we developed allow the monitoring framework to coalesce the events collected at various layers of abstraction, and to map these events back to the offending processes. We extended this mechanism to also provide the ability to block network events that attempt to exfiltrate data over a network connection. The mappings we infer are recorded in an audit trail, and we provide several mechanisms that help with data forensics efforts; for example, allowing an analyst to quickly reconstruct detailed information about what happened when such information is needed the most (e.g., after a system compromise). To demonstrate the practical utility of our framework, we show how our approach can be used to glean insightful information on behavioral profiles of malware activity after a security breach has been detected.

## ACKNOWLEDGMENTS

This work is supported in part by the National Science Foundation under awards OCI-1127361 and CNS-0915364

## REFERENCES

- [1] N. Provos, D. McNamee, P. Mavrommatis, K. Wang, and N. Modadugu, "The Ghost in the Browser: Analysis of Web-based Malware," in *First Workshop on Hot Topics in Understanding Botnets*, 2006.
- [2] J. Franklin, A. Perrig, V. Paxson, and S. Savage, "An Inquiry into the Nature and Causes of the Wealth of Internet Miscreants," in *14<sup>th</sup> ACM conference on Computer and communications security*, 2007, pp. 375–388.
- [3] D. Farmer and W. Venema, *Forensic Discovery*. Addison-Wesley, 2006.
- [4] R. Goldberg, "Survey of Virtual Machine Research," *IEEE Computer Magazine*, vol. 7, no. 6, pp. 34–35, 1974.
- [5] P. Chen and B. Noble, "When Virtual is Better than Real," in *Workshop on Hot Topics in Operating Systems*, May. 2001, pp. 133–138.
- [6] Sean Peiset and Matt Bishop and Keith Marzullo, "Computer Forensics in Forensics," *ACM Operating System Review*, vol. 42, 2008.
- [7] G. H. Kim and E. H. Spafford, "The Design and Implementation of Tripwire: a File System Integrity Checker," in *2<sup>nd</sup> ACM Conference on Computer and Communications Security*, 1994, pp. 18–29.
- [8] D. Vincenzetti and M. Cotrozzi, "ATP - Anti Tampering Program," in *USENIX Security*, 1993, pp. 79–90.
- [9] A. Goel, K. Po, K. Farhadi, Z. Li, and E. de Lara, "The Taser Intrusion Detection System," in *Symposium on Operating Systems Principles*, Oct. 2005.
- [10] S. Jain, F. Shafique, V. Djeri, and A. Goel, "Application-Level Isolation and Recovery with Solitude," in *EuroSys*, Apr. 2008, pp. 95–107.
- [11] K. Muniswamy-Reddy, D. Holland, U. Braun, and M. Seltzer, "Provenance-aware Storage Systems," in *USENIX Annual Technical Conference*, 2006, pp. 43–56.
- [12] K.-K. Muniswamy-Reddy, P. Macko, and M. Seltzer, "Provenance for the Cloud," in *USENIX Conference on File and Storage Technologies (FAST)*. Berkeley, CA, USA: USENIX Association, 2010.
- [13] T. Garfinkel, B. Pfaff, J. Chow, M. Rosenblum, and D. Boneh, "Terra: A Virtual Machine-Based Platform for Trusted Computing," in *ACM Symposium on Operating System Principles*, 2003, pp. 193–206.
- [14] S. T. Jones, A. C. Arpaci-Dusseau, and R. H. Arpaci-Dusseau, "Antfarm: Tracking Processes in a Virtual Machine Environment," in *USENIX Annual Technical Conference*, 2006.
- [15] S. Jones, A. Arpaci-Dusseau, and R. Arpaci-Dusseau, "Geiger: monitoring the buffer cache in a virtual machine environment," *12th International conference on Architectural Support for Programming Languages and Operating Systems(ASPLOS)*, vol. 41, no. 11, pp. 14–24, 2006.

- [16] X. Jiang, X. Wang, and D. Xu, "Stealthy Malware Detection through VMM-based "out-of-the-box" Semantic View Reconstruction," in *14<sup>th</sup> ACM conference on Computer and Communications Security*, 2007, pp. 128–138.
- [17] S. King and P. Chen, "Backtracking Intrusions," *19<sup>th</sup> ACM Symposium on Operating Systems Principles*, Dec 2003.
- [18] S. T. King, Z. M. Mao, D. G. Lucchetti, and P. M. Chen, "Enriching intrusion alerts through multi-host causality," in *Network and Distributed System Security Symposium*, 2005.
- [19] L. Litty, H. Lagar-Cavilla, and D. Lie, "Hypervisor Support for Identifying Covertly Executing Binaries," in *USENIX Security Symposium*, Aug. 2008, pp. 243–257.
- [20] B. D. Payne, M. Carbone, and W. Lee, "Secure and flexible monitoring of virtual machines," *Annual Computer Security Applications Conference*, pp. 385–397, 2007.
- [21] S. Krishnan, K. Z. Snow, and F. Monrose, "Trail of bytes: Efficient Support for Forensic Analysis," in *17<sup>th</sup> ACM conference on Computer and Communications Security*, Oct 2010.
- [22] F. Buchholz and E. Spafford, "On the Role of File System Metadata in Digital Forensics," *Digital Investigation*, vol. 1, no. 4, pp. 298 – 309, 2004.
- [23] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, "Xen and the Art of Virtualization," in *19<sup>th</sup> ACM Symposium on Operating Systems Principles*, 2003, pp. 164–177.
- [24] F. Leung, G. Neiger, D. Rodgers, A. Santoni, and R. Uhlig, "Intel Virtualization Technology: Hardware Support for Efficient Processor Virtualization," *Intel Technology Journal*, vol. 10, 2006.
- [25] A. Dinaburg, P. Royal, M. Sharif, and W. Lee, "Ether: Malware Analysis via Hardware Virtualization Extensions," in *15<sup>th</sup> ACM Conference on Computer and Communications Security*, 2008, pp. 51–62.
- [26] D. E. Denning and P. J. Denning, "Certification of Programs for Secure Information Flow," *Communications of the ACM*, vol. 20, no. 7, pp. 504–513, 1977.
- [27] S. Chen, J. Xu, N. Nakka, Z. Kalbarczyk, and R. K. Iyer, "Defeating Memory Corruption Attacks via Pointer Taintedness Detection," in *IEEE International Conference on Dependable Systems and Networks (DSN)*, 2005, pp. 378–387.
- [28] A. Slowinska and H. Bos, "Pointless Tainting? Evaluating the Practicality of Pointer Tainting," in *EuroSys*, Apr. 2009.
- [29] S. Quinlan and S. Dorward, "Venti: A New Approach to Archival Data Storage," in *USENIX Conference on File and Storage Technologies*, 2002, pp. 89–101.
- [30] A. W. Leung, S. Pasupathy, G. Goodson, and E. L. Miller, "Measurement and Analysis of Large-scale Network File System Workloads," in *USENIX Annual Technical Conference*, 2008, pp. 213–226.
- [31] A. Tirumala, F. Qin, J. Dugan, J. Ferguson, and K. Gibbs, "Iperf: The tcp/udp bandwidth measurement tool," *URL: http://dast.nlanr.net/Projects/Iperf*, 2004.
- [32] NIST, "National Software Reference Library," 2009.
- [33] F-Secure, "MBR Rootkit, A New Breed of Malware," See <http://www.f-secure.com/weblog/archives/00001393.html>, 2008.
- [34] T. Garfinkel, K. Adams, A. Warfield, and J. Franklin, "Compatibility is not Transparency: VMM Detection Myths and Realities," in *11<sup>th</sup> USENIX workshop on Hot topics in operating systems*, 2007, pp. 1–6.
- [35] X. Chen, J. Andersen, Z. Mao, M. Bailey, and J. Nazario, "Towards an Understanding of Anti-virtualization and Anti-debugging Behavior in Modern Malware," in *Dependable Systems and Networks*, June 2008, pp. 177–186.
- [36] T. Garfinkel and M. Rosenblum, "A Virtual Machine Introspection Based Architecture for Intrusion Detection," in *Network and Distributed Systems Security Symposium*, 2003, pp. 191–206.
- [37] L. Cavallaro, P. Saxena, and R. Sekar, "On the limits of information flow techniques for malware analysis and containment," *Detection of Intrusions and Malware, and Vulnerability Assessment*, pp. 143–163, 2008.



**Srinivas Krishnan** Srinivas Krishnan is a Ph.D. candidate at the University of North Carolina at Chapel Hill under Dr. Fabian Monrose. He obtained his Master's degree in Computer Science from University of North Carolina at Chapel Hill in 2007. After graduating, he worked for VMware in their performance R&D division, where he headed efforts on developing performance benchmarks for virtual desktop infrastructure. He began his PhD in 2008 and spent the summer of 2010 at Microsoft Research working on datacenter routing. His primary interests

are in computer and network security with a focus on developing mechanisms to monitor and leverage DNS infrastructure for detecting network exploits. He is also keenly interested in developing virtualization methods for providing monitoring and forensic capabilities.



**Kevin Z. Snow** Kevin Z. Snow is a Ph.D. candidate at University of North Carolina at Chapel Hill under the advisement of Dr. Fabian Monrose. He obtained dual Master's degrees in Computer Science and Information Security from Johns Hopkins University in 2007. After graduating, he joined the Cyber Operations group at Johns Hopkins Applied Physics Laboratory, primarily working in the areas of information system attack and scalable network monitoring. He began pursuit of his Ph.D. in 2009 and spent the summers of 2011 and 2012 with

Google's Safe Browsing Team. His primary interests are in computer and network security with a particular focus on practical methods of software exploitation, scalable detection of information system attacks, and attack forensic techniques that add value for network operators.



**Fabian Monrose** Dr. Fabian Monrose is an Associate Professor of Computer Science at University of North Carolina at Chapel Hill. He was awarded a Ph.D. in Computer Science from the Courant Institute of Mathematical Sciences, New York University, in 1999. After graduating, he joined the Secure Systems group at Bell Labs, Lucent Technologies, where he remained until 2002. From 2002 through 2008 he held appointments in the department of Computer Science at Johns Hopkins University and the Information Security Institute. His interests span

the fields of Networking and Security and includes such diverse topics as Traffic Classification, Computer Forensics, User Authentication, and Privacy. In 2006, he received a CAREER award from the National Science Foundation. He has published over 70 peer reviewed papers in flagship security conferences and journals, and has been the recipient of several best paper awards.