

APTSHIELD: A Stable, Efficient and Real-time APT Detection System for Linux Hosts

Tiantian Zhu, Jinkai Yu, Tieming Chen*, Jiayu Wang, Jie Ying, Ye Tian, Mingqi Lv, Yan Chen, Yuan Fan and Ting Wang

Abstract

Advanced Persistent Threat (APT) attack usually refers to the form of long-term, covert and sustained attack on specific targets, with an adversary using advanced attack techniques to destroy the key facilities of an organization. APT attacks have caused serious security threats and massive financial loss worldwide. Academics and industry thereby have proposed a series of solutions to detect APT attacks, such as dynamic/static code analysis, traffic detection, sandbox technology, endpoint detection and response (EDR), etc. However, existing defenses are failed to accurately and effectively defend against the current APT attacks that exhibit strong persistent, stealthy, diverse and dynamic characteristics due to the weak data source integrity, large data processing overhead and poor real-time performance in the process of real-world scenarios.

To overcome these difficulties, in this paper we propose APTSHIELD, a stable, efficient and real-time APT detection system for Linux hosts. In the aspect of data collection, audit is selected to stably collect kernel data of the operating system so as to carry out a complete portrait of the attack based on comprehensive analysis and comparison of existing logging tools; In the aspect of data processing, redundant semantics skipping and non-viable node pruning are adopted to reduce the amount of data, so as to reduce the overhead of the detection system; In the aspect of attack detection, an APT attack detection framework based on ATT&CK model is designed to carry out real-time attack response and alarm through the transfer and aggregation of labels. Experimental results on both laboratory and Darpa Engagement show that our system can effectively detect web vulnerability attacks, file-less attacks and remote access trojan attacks, and has a low false positive rate, which adds far more value than the existing frontier work.

1. Introduction

Advanced Persistent Threat Attacks are escalating to the harm of the current society, and it is often organized by

groups of hackers with certain national, governmental or other organizational backgrounds—the hackers are usually well-organized, well-targeted, highly skilled and aggressive, often against the government, core infrastructure (e.g., energy, transportation, communication) and key industries (e.g., military, finance, health care). APT attacks can pose a huge security threat, including confidential data leakage and system integrity damage. APT attacks have occurred frequently in recent years, showing a high incidence of high risk in the world, such as: Stuxnet worm attack in Iranian nuclear power plant [1], BlackEnergy virus attack in Ukraine’s power system [2], and the user information leakage attack in Target [3]. The COVID-19-themed attacks have also been frequently reported [4], e.g., APT groups have tried to attack firms working on COVID-19 vaccines and they have used spear-phishing emails to entice users to download and execute malicious attachments so as to steal target-related data and destroy medical infrastructure [5].

APT detection has become an important research topic widely concerned with academia and industry field. A series of traditional attack detection schemes such as static code analysis [6, 7], dynamic sandbox detection [8, 9], malicious traffic analysis [10, 11], and hooking technology [12, 13] can be applied to fight against APT attacks. However, the detection effect of the above methods is not ideal due to the problems such as weak de-obfuscate ability, high computational overhead, and low system stability in the practical use. Furthermore, with the extension of the network boundaries and the increment of various 0-day vulnerabilities, it is almost impossible to continue to use these traditional methods to detect APT attacks. Therefore, the idea of Detection and Response (DR) came into being, which is divided into Endpoint Detection and Response (EDR) and Network Detection and Response (NDR). In EDR, the relevant security agent will collect and analyze activity data of the application program when it is running in the user’s host/endpoint [14]. While NDR will monitor how threats enter the network and how they move laterally in the network [15]. Both EDR and NDR can automatically respond to identified threats to remove or contain them.

Nowadays, the market for EDR solutions is expanding

rapidly to meet the urgent need for more efficient endpoint protection and potential vulnerability detection. To this end, the DARPA Transparent Computing (TC) program has tried to organize multiple scientific research units to conduct engagements to enable the prompt detection of APTs and other cyber threats [16]. EDR tools usually record a large number of system events to a central database. By adopting techniques such as indicators of compromise (IOC), behavior analysis, and machine learning to analyze data, it is hoped that threats will be detected and responded to at an early stage.

However, in the life cycle of APT attacks, the current EDR system (especially the EDR system of the Linux host) does not form a comprehensive and effective solution to the problems of selecting data sources for detection, massive data analysis and storage, association of suspicious behaviors in context, detection of 0-day threats, as well as energy consumption and real-time performance in detecting APT attacks. The main challenges of the current EDR system based on the Linux host are as follows:

(1) How to select reliable, stable and semantically rich data sources. To carry out complete and accurate APT detection and forensics analysis in the enterprise environment, it is necessary to retain the complete log data of the entire enterprise for subsequent analysis. There are three main log collection methods in the existing attack detection work: The first is the *network data flow*, analysts only need to use a switch that supports port forwarding and a server that analyzes network traffic to obtain relevant data. However, encryption technologies are widely used by sophisticated attackers, making the effective information less available in network data; The second is the *application log*, in this case, both instrumentation and collecting logs directly from the application are limited to the application itself. Analysts cannot observe a global view of the whole system. In other words, the analyst cannot associate the information of application level with that of system level. The third is the *endpoint system log*. Nowadays, there are a few studies focusing on the collection and analysis of system events. However, the existing system-based log collection tools are varied, and few work are able to analyze and evaluate these tools under the current Linux system by combing with collecting performance, data integrity, availability and other comprehensive factors.

(2) How to reduce the amount of data required for real-time detection to improve the detection efficiency. The duration of APT attacks is usually much longer than other attacks. According to the report from Trustwave [17], the average latency time of APT attacks is about 83 days, and some of them are as long as several years. Massive data of APT attacks have made new requirements for analysts: It is critical to improve real-time detection efficiency and reduce storage overhead by effective data compaction strat-

egy. Some existing data compaction work [18–23] tries to use fine-grained taint tracking technology to delete redundant events. However, these methods always rely on known software models which are lack of generalization. Although some studies [24–26] have proposed general data compaction methods based on audit log, their algorithms read the dependency graph composed of long-term log data into memory at one time, resulting in huge computing and memory overhead. By adopting the above methods [24–26] in the real-world scenario, data compaction and attack detection cannot be guaranteed in a real-time manner.

(3) How to construct a real-time APT detection framework with high adaptability, high coverage, high precision, low false positives and constant memory overhead.

On the one hand, traditional single point intrusion detection methods (e.g., network traffic analysis, software static feature detection, dynamic sandbox detection and hooking technology) cannot have the ability of anti-escape, reliable detection and real-time alarm at the same time under the threat scenarios with the characteristics of persistence, concealment and diversity of APT attacks. On the other hand, the contextual data analysis studies [27,28] usually preserve context information via the provenance graph. However, the size of the provenance graph will explode over time due to the long duration of APT attacks, rendering these approaches inevitably suffer from efficiency and memory problems [29]. Furthermore, various new attack channels and carriers emerge one after another in APT, such as web vulnerability attacks, file-less attacks (also known as in memory attacks or living-off-the-land attacks) and remote access trojan attacks, which are difficult to detect in the initial stage of intrusion. The ideal detection scheme should have good expansibility, be able to cover different types of sophisticated APT attacks, effectively define APT related suspicious behaviors, and comprehensively utilize the full contextual attack chain to realize real-time and accurate detection of attacks.

To cope with the above challenges, we propose a stable, efficient and real-time APT detection system for Linux hosts, called APTSHIELD¹. Firstly, in order to select reliable, stable and semantically rich data sources, we make a comprehensive and detailed evaluation on Linux system log collection tools. Secondly, for the purpose of reducing the storage overhead and improve the detection efficiency, we reduce the amount of log data by means of redundant semantic skipping and non-viable node pruning. Finally, we construct an APT attack detection framework based on ATT&CK model [30], and the framework can carry out real-time attack response and alarm through the transfer and aggregation of labels. In general, the contributions of this paper are as follows:

¹APTSHIELD for Stable, Efficient and real-time APT Detection system for Linux Hosts.

- Different from previous studies, we make a detailed comparison and analysis of the advantages and disadvantages of the existing data source. We deploy multiple indicators and then obtain the optimal data sources for stable APT detection on Linux through performance analysis and comprehensive judgment.
- We use redundant semantics skipping and non-viable entity pruning to improve the efficiency of real-time APT detection and reduce data storage overhead for forensic analysis. Our data compaction methods can be carried out in real-time streaming data to ensure the performance of the detection system. Also, the compaction effect is better than that of existing studies without affecting the accuracy of final results.
- Based on the ATT&CK model, we constructed an information aggregation framework through system data flow and control flow. With the help of Tactics, Techniques and Procedures (TTP), the atomic suspicious characteristics of system entities and their transmission rules are defined in the framework to realize the aggregation of the contextual information of APT attacks. Different from traditional single-point detection methods, our framework can aggregate the information of the attack chain into specific entities, and realize the whole chain detection and real-time alarm of APT attacks with constant memory overhead.
- We implement a stable, efficient and real-time APT detection prototype system on Linux and conduct experiments on the dataset of Darpa Engagement as well as the dataset that simulates APT attacks in real-world scenarios. The experimental results show that our system can effectively detect web vulnerability attacks, file-less attacks and remote access trojan attacks in a real-time manner, and has a low false positive rate, which adds far more value than the existing frontier studies.

The remainder of this article is organized as follows. Section 2 surveys the relevant work from three aspects: data collection, data reduction and APT detection. Section 3 describes the preliminary background for our work. System design and evaluation are described in Section 4 and Section 5. We conclude and discuss the limitations and improvements of APTSHIELD in Section 6.

2. Related Work

In this section, we review notable studies in APT detection and compare with them in data collection, data reduction and detection performance to highlight the novelty of our approaches.

Data collection. In order to carry out practical and effective APT detection, it's crucial to select the suitable data source which meets the characteristics of non-tampering, low resource consumption, and stability. Network data flow [31, 32] are widely used by previous studies for intrusion analysis and detection, but its universality was insufficient to record all attack operations. Furthermore, encryption technology makes it difficult for analysts to obtain effective information. The application data [33] shows the inherent characteristics of the program when it is running, but the analyst cannot associate the information of application level with that of system level. Also, some studies [34, 35] use instrumentation to collect taint records, which may cause huge memory overhead. Some logging tools are also used by researchers, such as NanoLog [36], Log4j2 [37], Spdlog [38], Glog [39], and BoostLog [40]. Nanolog is a nanosecond logging system, which is 1-2 orders of magnitude faster than other systems, and its throughput is capable of reaching 80 million log messages per second. However, Nanolog is implemented in C++ and it can only operate on some static strings. Although the remaining systems (i.e., Log4j2, Spdlog, Glog, and BoostLog) are able to meet some requirements, such as low resource consumption and fidelity, they are highly coupled with the application itself, and cannot be used as a general log infrastructure. Also, the above systems cannot log kernel events to monitor all user actions. In view of the deficiencies of the above data sources, kernel-based data collection tools are used by researchers, such as Auditd [41], Sysdig [42], Ltng [44], and Auditbeat [45]. However, there are few studies that conduct a complete and detailed analysis of these popular data collection tools based on the needs of the real-world environment.

Different from the previous work, in this paper we compare the performance of the above four kernel-based data collection tools when the system is no/full load through experiments. The details will be discussed in Section 4.3. Sysdig is able to depict complete kernel data and parameters, but its requirements for CPU and memory are very high due to its excessive content. The overall performance of Auditd and Ltng is better than Sysdig, but the output of Ltng is the binary stream, and it takes a lot of time to recover the semantic related data. Auditd has many data filtering modules, it can reduce resource consumption through custom filtering strategies, thereby meeting all the requirements of the detection system for data sources.

Data reduction. Due to the huge amount of data required (the average latency time of APT attacks is about 83 days) to detect APT attacks, it is critical to improve real-time detection efficiency and reduce storage overhead by effective data compaction strategy. To reduce the system events without affecting the APT detection results is a hot research direction in recent years.

LogGC [18] pioneered the idea of garbage collection for

audit log. The authors of LogGC combined with BEEP [19] to delete events that had no lasting impact on the system. Subsequently, NodeMerge [46] presented a template-based data compaction system with the assistance of an improved FP-Growth algorithm. Some specific patterns such as image loading and system configuration were extracted for data compaction. Similarly, Conan [29] prefiltered the duplicated read events through semantic recognition to reduce the detection efficiency. Zhu et al [59] maintained a long list to record redundant events. However, the above three methods are only applicable to specific event types, and have limitations in complex scenarios (e.g., when files are accessed by multiple processes, or there are a large number of file write operations, or the real-time performance of the system will deteriorate as the collection time increases).

ProTracer [20] tried to improve the compaction rate by dividing the target program into multiple units to perform fine-grained taint tracking. Later, MPI [21] improved ProTracer by utilizing a semantics aware program annotation and instrumentation technique to partition execution, which was able to generate execution partitions with rich semantic information. To avoid application instrumentation or kernel modification, MCI [22] utilized LDX [23] to acquire precise causal models for a set of primitive operations. The compaction effect of above methods depends on a large number of software models. However, there are a lot of unpredictable software running in the real-world production environment, and it is very difficult to ensure the coverage of these methods. In addition, software updates may invalidate the above methods.

Xu et al. [47] presented the concept of trackability. By aggregating events under the same trackability (i.e., deleting multiple equivalent data streams and retaining only one of them), they could reduce a large portion of data while preserving events relevant to a forensic analysis. However, this method only considers the characteristics of a single node instead of the global semantics, which makes the compaction effect very limited. To solve this problem, both Hossain et al. [25] and Hassan et al. [26] presented the data compaction algorithm based on the global semantics of the provenance graph. But during the calculation, the whole data of the provenance graph needs to be read into memory at once, which brings additional I/O and memory overhead. Furthermore, the above methods cannot process real-time streaming data and guarantee the real-time performance of APT detection.

In contrast, APTSHIELD is able to compact (i.e., redundant semantics skipping and non-viable node pruning) the real-time streaming data to improve the efficiency of real-time APT detection, reduce data storage overhead for forensic analysis, and achieve a better compaction effect than existing work without affecting the accuracy of final result.

APT detection. APT attackers usually exploit zero-day vulnerabilities to compromise the target and conduct long-term penetration. Traditional attack detection schemes such as static code analysis [6, 7], dynamic sandbox detection [8, 9], malicious traffic analysis [10, 11], and hooking technology [12, 13] are not ideal when facing with the characteristics of persistence, stealth, diversity, and clear goals in APT attacks. In order to effectively monitor APT attacks, endpoint detection and response (EDR) is gradually used by researchers. The provenance graph is an ideal method for threat modelling with the ability of powerful semantic expression and attacks historic correlation in the EDR system [48]. Provenance graph generally records the coarse-grained data from the kernel level, including subjects (e.g., processes and sockets), objects (e.g., files), and edges (e.g., system events).

BackTracker [51] and PriorTracker [52] tried to use provenance graph to identify the entry point and the effect of the attack through backward tracking and forward tracking, respectively. Sun et al. [49] utilized Bayesian Networks to identify zero-day attack paths on the provenance graph. Similarly, NoDoze [53] performed attack triage within the provenance graph to find anomalous paths. During the provenance-based attack investigation, analysts usually came across the dependence explosion problem (i.e., an output event is assumed to be causally dependent on all preceding input events, and an input event is assumed to have causal influence on all subsequent output events). To mitigate this problem, Ma et al. [20] presented a lightweight provenance tracing system to reduce the memory overhead via unit-based execution partitioning [19], as well as their improved studies [21, 54]. As discussed in data reduction, the above three methods mainly rely on software models and are not universal.

Due to the persistence of APT attacks, it is difficult for a security analyst to pick out “needle-in-a-haystack” attacks. *Learn-based APT detection* on long-term provenance graphs has been proposed by some researchers. Barre et al. [56] tried to extract statistic features of key processes, by adopting a random forest model, their system was able to deliver a detection rate of about 50%. The results show that simple feature engineering without considering the context of the provenance graph cannot effectively characterize complex APT attacks. Berrada et al. [57] designed a comprehensive experiment on provenance graphs by combining existing techniques for APT anomaly detection and they found that simple score or rank aggregation techniques were effective at improving detection performance. Han et al. [50] presented UNICORN which could model system behaviors via structured provenance graphs with a graph sketching technique. Although UNICORN requires no prior expert knowledge of APT attack patterns or behaviors, the anomaly-based system will undoubtedly introduce a large

number of false alarms, which is difficult to be practical in real-world scenarios.

In fact, *rule-based APT detection* is more in line with commercial needs, it has been shown that rule-based EDR systems are suitable for addressing the noise problem [55] (i.e., false alarms and duplicate alerts). Expert policies have been proposed by Sleuth [27] and Holmes [28] for attack reconstruction on provenance graphs. To match possible exploits of localized components in the provenance graph, empirical labels and prior specifications were used by Sleuth and Holmes, respectively. However, the size of the provenance graph would explode over time due to the long duration of APT attacks, rendering these two approaches inevitably suffer from efficiency and memory problems. Poirot [58] tried to detect APT attacks by measuring the offline similarity between a provenance graph and a query graph. The query graph was obtained based on the expert knowledge from cyber threat intelligence (CTI). The disadvantage of Poirot is that only the occurred attacks are able to be further detected. It is a critical limitation given that composing the elaborate description of a new category of APT requires significant forensics efforts. Conan [29] detected APT attacks by a three-phase model and pre-defined rules. Although Conan is capable of capturing the suspicious behaviour in a real-time manner, it has the following two problems. First, in the case of a large number of concurrent operations in the system (e.g., a large number of file read and write operations at the same time), the real-time performance of Conan cannot be guaranteed because the event generation rate of the target host may be much larger than the event throughput of the detection system. Second, the three-phase model proposed by Conan only reflects the path of suspicious data/control flow, and it has no performance on the tactics used by attackers. That is to say, it lacks interpretable guidance for the analysis of security analytics.

Unlike the previous methods, APTSHIELD can detect long-term APT attacks with a constant level of memory overhead. By adopting the data compaction strategies, APTSHIELD is capable of guaranteeing the real-time performance. Also, APTSHIELD is able to cover different types of sophisticated APT attacks (i.e., web vulnerability attacks, file-less attacks and remote access trojan attacks), effectively define APT related suspicious behaviors based on ATT&CK model, and comprehensively utilize the full contextual attack chain to realize accurate detection of attacks.

3. Background

In this section, we first introduce system entities and system events, followed by the dependency graph composed of them. Then, we describe the dependency relationship in the dependency graph. Finally, we introduce how to use prove-

nance graph for forensic analysis.

3.1. System Entity and System Event

Generally, the system entity is able to be divided into two categories: subject and object. The subject is the initiator of the event. In most cases, the subject refers to the process. Specifically, users are treated as subjects when there is a user-related event. The object is the target of the event, such as files, network IPs, and channels. It is worth mentioning that processes (e.g., child processes) can also be treated as objects. System events usually record the actions initiated by the subject to the object, such as file reading, file writing, process creation, etc.

3.2. Dependence Graph and Dependency

In order to facilitate efficient analysis of system logs, system entities and events are usually displayed with a dependence graph [51] visually. The dependence graph is also known as a provenance graph or a dependence graph. In these graphs, the nodes including subject and object, and the edges represent the event relationship between subject and object. Dependency graph is a directed graph, and the direction of the event is from subject to object. For each edge, there will be a timestamp and an event name to represent the event sequence and the meaning of the event, respectively. In some research studies of data compaction [18, 29, 59], system events are able to be merged with multiple timestamps or time ranges.

Dependency graph G can be represented as a combination of (V, E) , where V represents all nodes (system entities) and E stands for all edges (system events). For any edge $e \in E$, there is $e = (u, v, t)$, where $u, v \in V$, u represents the subject, V represents the object, and t represents the time stamp of the event. For the two edges in the dependency graph, $e1 = (u1, v1, t1)$, $e2 = (u2, v2, t2)$, we consider that there is a dependence between $e1$ and $e2$ if $v1 = u2$ and $t1 < t2$.

3.3. Forensic Analysis

Forensic analysis is also known as trace analysis, and its main purpose is to help analysts understand when, how and what impact the attack has been made. In the process of APT analysis, analysts have two tasks. The first one is to determine the entrance of the attack (e.g., the initial process and file entering the terminal and the source IP/port of the attack). The analysis process to achieve this goal is called backward analysis, which is tracing suspicious information flow from the process marked as suspicious through the opposite direction of information flow. Backward analysis was first proposed by BackTracker [51]. During backward analysis, the timestamp is used to judge the causal relationship between different events. For the second task, forward analysis is used to analyze the impact of attacks [25, 59], such as

accessing to sensitive information, tampering with system configuration, etc. The origin of forward analysis is usually the attack entry point obtained by backward analysis or any point in the attack chain.

4. Methods

In this section, we first introduce the threat model and the overall architecture of our proposed APT detection system. We then discuss several important topics in its design, including data collection, data compaction, and APT detection framework.

4.1. Threat Model

Based on famous APT attack, RSA SecurID tokens leakage [60], we consider the following attack scenarios: an employee received an email with the words "recruitment plan". The employee downloaded and opened the attachment, then was hit by the latest 0-day vulnerability. At the same time, the attacker established a command&control connection with the employee's host, continuously downloaded and executed the remote access Trojan that the attacker had prepared, which was not able to be detected via the anti-virus software. As a result, the attacker now had access to the target host in the corporate environment. Further, the attacker used this machine as a springboard to perform lateral movement. And the other employee's or senior manager's host that associated with the employee's were successively controlled by the attacker. Then the attacker could stay for months or even years until he completed the ultimate goal, such as the theft of confidential documents or the destruction of the integrity of the system.

In the threat model of this article, the logs generated by the operating system are considered credible. We assume that the system is not attacked before the start of data collection. Attackers can compromise the target host in a variety of ways, including but not limited to 0-day vulnerabilities, social engineering, virus-carrying USB flash drive, etc. The primary task of the attacker is to run the malicious code on the victim's machine, and then collect target information through remote command&control. Attackers will hide themselves in normal system activities through disguise, with the ultimate goal of stealing high-value data or destroying the integrity and availability of the system.

4.2. System Overview

The architecture of APTSHIELD is illustrated in Figure 1. It consists of three parts: data collection module, data compaction module, and APT detection framework. Firstly, kernel-based tools are adopted for data collection. Through comprehensive analysis, we choose auditd as the collector on the client side since it provides sufficient data with low overhead. The collected system logs (including entities and events) will construct a dependence graph. Secondly,

redundant semantics skipping and nonviable node pruning will be adopted to compact system logs. The main advantage of data compaction is to improve the efficiency of real-time APT detection and reduce data storage overhead for forensic analysis. Thirdly, APT detection framework is proposed based on the ATT&CK model. The atomic suspicious characteristics of system entities (i.e., process labels and file labels) and their transfer rules (through events with different semantics) are defined in the framework to realize the aggregation of the contextual information of APT attacks. The APT will be alerted from a specific entity (through judgement rules), and finally the related attack chain will be found through forensic analysis.

4.3. Data Collection

In order to detect APT in a real-time manner with high efficiency, it is crucial to select optimal data sources and data collectors. To this end, in this section we first sort out the data sources for APT detection under the Linux system, and then analyze and select the best data collector that meets the requirements.

4.3.1 Data Requirements

According to the ATT&CK model, the host-based APT detection system needs to collect relevant data from different stages (e.g., initial access, untrusted execution, data exfiltration, etc.). We summarize requirements of both data collector and data sources into the following three aspects:

(1) Tamper-proof. APT attacks are camouflaged. In order to avoid detection, APT attackers often hide their attack traces (e.g., bash history and log message can be easily modified/deleted by an attacker, thereby increasing the difficulty of attack detection). We need to select credible data sources to ensure that intruders cannot avoid detection by tampering with data.

(2) Stability and Low Overhead. APT attacks are persistent. That means data collection is a long-term process. Firstly, the data collector needs to be directly deployed on the user's host. If the data collection takes up too much resources, it will affect the user's normal usage. Secondly, the detection system needs to perform stable and real-time detection and alarm, this requires data to be collected and transmitted in real-time without any data loss. Thirdly, the operation of the data collector cannot affect the stability of the system (e.g., it should not be requested to modify the kernel)

(3) Rich semantics. APT attacks are diverse. It is necessary to analyze a series of data such as file IO data, network IO data, inter-process communication data, process attribute modification data, sensitive system operation data, and file attribute modification data.

Based on three requirements above, the kernel-level logs

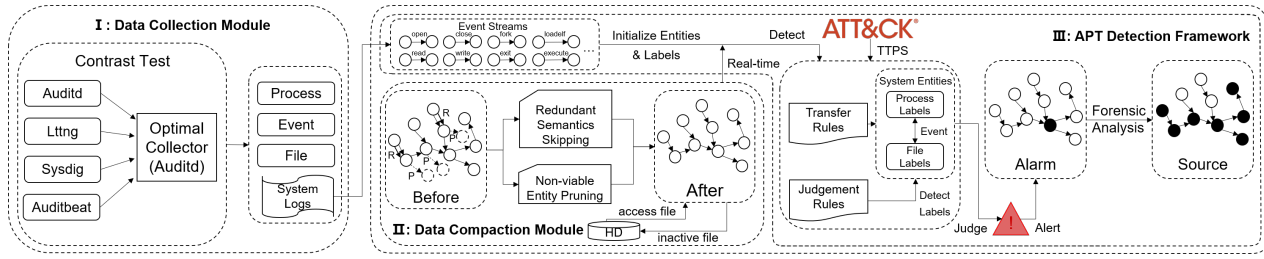


Figure 1. System Overview of APTSHIELD. Data collection module, data compaction module and APT detection framework are the key components of APTSHIELD. The kernel-based tool (Auditd) is deployed on the client side for system logs collection. By adopting redundant semantics skipping and nonviable node pruning, data compaction module will help data storage overhead. APT detection framework is responsibility for real-time aggregation of the contextual information of APT attacks with the help of ATT&CK model. Once the threat/attack occurs, APTSHIELD can find the entire attack chain through forensic analysis.

Table 1. Performance comparison of four different kernel-based data collectors. We tested the CPU (single core) and memory usage of each collectors on full-load and no-load. All the results are the average based on 10 tests independently.

Collector	No-load		Full-load	
	CPU (%)	Memory (MB)	CPU (%)	Memory (MB)
Auditd	0.2	31.2	2.0	31.2
Lttng	2.0	403.6	23.0	467.6
Sysdig	19.0	108.4	22.0	463.4
Auditbeat	1.2	136.4	7.1	170.4

are used as the main data source. We analyze and select the best data collector currently available on Linux, which will be described in details in the next section.

4.3.2 Data Source Analysis and Selection

As introduced in Section 2, current data collectors that are active in the community are sysdig, auditd, lttng, and auditbeat. With the Intel Xeon E5 CPU (8 cores) and 64G memory running on Ubuntu 16.04, we analyze and compare the performance (CPU and memory usage) of these data collectors when the host is no-load and full-load (i.e., use commands to increase the CPU usage to more than 95%. The executed commands include but are not limited to file reading and writing, file downloading, inter-process communication, executing shell scripts, accessing pages, etc.), respectively. The results are shown in Table 1.

It can be seen from Table 1 that auditd and auditbeat have better overall performance when the server is fully loaded. The overhead (memory) of auditd is the same under both no-load and full-load because its memory usage is related to the buffer size. In the actual test, the buffer size can be controlled within 40MB to meet the data demand under full load. If sysdig is used as a collector, there will be service

crashes and data loss when the system is fully loaded. The output data generated by lttng is a binary byte stream, and parsing requires a lot of costs. Considering that auditd is a module that comes with the kernel and has the lowest cost, we finally select auditd as the data collector in this article. In addition, we use the stress testing tool UnixBench [61] to monitor the state of auditd when it is turned on/off. The results show that when auditd is turned on, the total score of UnixBench drops about 0.3% compared to that is turned off, which is basically negligible.

4.4. Data Compaction

If the collected raw data is sent directly to the server for processing, it will consume a lot of storage and network resources due to the long duration of the APT attack. Therefore, in the real-world enterprise environment, these data should be compacted first, and then sent to the storage/detection server. The data compaction method should meet the following conditions: (1) **Real-time**. The data consumption rate should be greater than the data generation rate to avoid the impact of cached data on the real-time performance of the detection system; (2) **High efficiency**. The data compaction method has low CPU and memory usage, and will not affect the operating system itself. Also, it will greatly reduce the data transmission bandwidth and data storage cost; (3) **High accuracy**. The compacted data should maintain the dependencies in the original data, and should not have a negative impact on the APT detection and forensic analysis of APT attacks. In this section, we will first give the insight for data compaction. Then we introduce two algorithms in our data compaction method: redundant semantics skipping and non-viable node pruning.

4.4.1 Insight for Data Compaction

In order to perform data compaction without affecting the detection results, it is important to analyze the contextual semantics of dependency graph. In the attack detection

based on the dependency graph, it is necessary to analyze the interaction behavior of the system entities by judging the different semantics of the system events (that is, the edges of the dependency graph), and then find the path related to the attack. From the perspective of semantic transfer, we assume that entity u has an event $e1$ directed to entity v at time $t1$. It is considered that entity u has an impact on entity v at time $t1$. Then at time $t2$ ($t2 > t1$), there is an event $e2$ pointing to entity w from entity v , so it is considered that entity v has an impact on entity w at time $t2$. Since the event $e2$ occurs after the event $e1$, it can be considered that the entity u has an impact on the entity w , which is transmitted through the entity v .

Take Figure 2 as an example. The circle represents the process, the rectangle represents the file, the diamond represents the IP/DNS, the arrow represents the event, and the number on the arrow represents the time point of the events (the smaller the number, the earlier the event occurs). Process P is affected by the event generated by $x.com$ at time $t1$, and P obtains semantic information from the network. After that, the system event with process P as the initial entity will have a direct impact on other entities (i.e., file A , file B , and process Q). However, considering that process P is only affected by the event generated by $x.com$ before time $t5$. It can be judged that from time $t1$ to time $t5$, other entities pointed by process P are only affected by the network information from $x.com$ to P . After the event at time $t5$ occurs, other entities pointed by process P are simultaneously affected by the network information from both $x.com$ to P and $y.com$ to P . According to observations above, we can see that the impact of the event on the entity is lasting. From the perspective of contextual semantics, when an entity receives an event containing external information, the entity will contain certain characteristics transmitted by the event. And these characteristics will affect the flow of information from this entity. Our main idea/insight of data compaction is derived from these observations: **If the semantics of source entity has not changed and the destination entity has received multiple same information streams (events) from the source entity, these events will have the same semantics, which can be combined to delete redundant semantics.**

The result of redundant semantics skipping is shown in Figure 2. The data compaction process is as follows:

- (1) After the time $t1$, process P already has the semantics from $x.com$, and the event at $t2$ do not introduce any new semantics and can be deleted.
- (2) Although during $t5$ and $t7$ an event (at $t6$) from P to Q occurred, it only affects the semantics of Q and the semantics of P remains unchanged. The events at time $t5$ and $t7$ have the same semantics, and we can delete the event at time $t7$.
- (3) We can find that events at $t4, t8$, and $t9$ have the same




Figure 2. An example about Data compaction. It contains the original provenance graph G and the compacted graph G' .

subject and object. The event at time $t5$ has changed the semantics of P , so the event at time $t8$ should be retained. Since the events at time $t8$ and time $t9$ have the same semantics, we can delete the event at time $t9$.

Note that in this example, we assume that the semantics of IP/DNS (i.e., external network) was not changed. But in the real-world scenario, the external network may be compromised by an attacker, thereby causing semantic changes. To address this problem, we design a time window T to retain the network receiving events at a low frequency, which is inspired by [25, 59]. For example, if $T = 50$, then after every 50 time intervals, the semantics of IP/DNS will be refreshed and retained again. We choose $T = 50$ in this paper.

4.4.2 Redundant Semantics Skipping

APTSHIELD maintains a table to store the latest semantic information of all processes. When the semantics of the incoming event is the same as that in the semantic table, the event can be considered as redundant and deleted. If there are a large number of repeated read and write operations, deleting redundant semantic events can greatly reduce execution time and memory overhead.

The algorithm of redundant semantics skipping is shown in Algorithm 1. The input is the real-time streaming data (events) and the output is whether to delete the event or not. In addition, we use LST to record the set of (entity, event) that have been occurred. When processing the real-time event, the algorithm will determine whether the subject S exists in LST . If it exists, then determine whether the event e is the same as stored event e' . If it is the same, the event will be deleted. We found that when the size of LST gets larger, the efficiency of this algorithm will become lower. Therefore, we propose to empty the list when the LST reaches a certain size. In the test, it is found that when the threshold is set to 5, APTSHIELD has better performance and efficiency. In addition, when the event e is a WRITE/RCV event, the event e' stored in LST related to O need to be deleted. The reason is that the semantics of the

object O will be changed due to the WRITE/RECV event.

Algorithm 1 Data Compaction of Redundant Semantics Skipping

Input:

- (1) Streaming data (events) in chronological order, each event e_i contains a subject S_i and a object O_i
- (2) The set of (entity, event) that have been occurred recently are denoted as Latest Semantic Table LST

Output:

Whether the event is skipped

Initialize:

The set of (entity, event) $LST = \emptyset$

```

1: for  $e_i \in$  streaming events do
2:   if  $S_i$  not exists in  $LST.Keys$  then
3:      $LST$  add  $(S_i, e_i)$ 
4:   else
5:     if  $e_i$  equals to  $e'_i$  then
6:       Skip the event
7:       Continue
8:     else
9:       Delete  $(S_i, e'_i)$  in  $LST$ 
10:      Add  $(S_i, e_i)$  to  $LST$ 
11:      if  $O_i$  exists in  $LST.Keys$  then
12:        Delete  $(O_i, e'_i)$  in  $LST$ 
13:      end if
14:    end if
15:  end if
16:  if  $LST.size \geq 5$  then
17:     $LST = \emptyset$ 
18:    Continue
19:  end if
20:  if  $e_i \in \{WRITE, RECV\}$  then
21:    for  $(S'_i, e'_i) \in LST$  do:
22:      if  $O_i$  exists in  $e'_i$  then
23:        Delete  $(S'_i, e'_i)$  in  $LST$ 
24:      end if
25:    end for
26:  end if
27: end for

```

4.4.3 Non-viable Entity Pruning

After filtering out redundant semantics, we maintain a node tree to contain the relationship of all processes for efficient forensic analysis. For all process entities that do not have a parent node (i.e., independent processes), we connect them by creating a virtual root process. In order to keep a stable memory overhead when the system runs for a long time, non-viable entity pruning is used to reduce the number of processes. As shown in Algorithm 2, if a process satisfies: (1) it executes the exit event, (2) it does not have potential harmful functionalities (PHF, it will be introduced in Section 4.5.2) and child nodes (When the child node has PHFs, the deletion of the parent node will affect the forensic anal-

ysis), this process will be pruned from the node tree.

Moreover, we move inactive file nodes (in this article, inactivity is defined as no change for more than 5 minutes, and the time interval can be changed according to actual needs) to the disk to further reduce the memory overhead. They will be took out from the disk in real-time when needed.

Algorithm 2 Data Compaction of Non-viable Entity Pruning

Input:

- Streaming data (only EXIT events) in chronological order. Each event e_i contains a subject P_i

Output:

Whether the process is pruned

Initialize:

For all processes are stored in a Process Tree PT

```

1: for  $e_i \in$  streaming events do:
2:   if  $e_i \in \{EXIT\}$  then
3:     if  $isleaf(P_i) \ \& \ \neg PHF(P_i)$  then
4:       prune  $P_i$ 
5:     end if
6:   end if
7: end for

```

4.5. APT Detection Framework

In this section, we first introduce the insight for label-based APT detection. Then we explain the definition of labels and label delivery rules based on ATT&CK model. Finally we give the attack judgment rules.

4.5.1 Insight for Label-based APT Detection

Due to the persistence of APT attacks, the detection methods based on the dependency graph suffer from efficiency and memory problems. To detect long-term APT attacks with a constant level of memory overhead, we develop various adaptable labels based on ATT&CK model, and transfer contextual semantics between entities through control flow and data flow. **Control flow** mainly refers to the process creation relationship. When the attack involves multiple processes, tracking the untrusted control flow can realize the association of multiple malicious processes. **Data flow** refers to the dependence of data content, which reflects the flow path of data. The data flow can be divided into untrusted data flow and high-value data flow. The former one reflects the flow path of external untrusted data, which is often the intrusion path of suspicious codes; the latter one reflects the escape path of high-value contents in the attack after being stolen. Untrusted control flow, untrusted data flow, and high-value data flow reflect the nature of the attack. By tracing the above three information flows, analysts can easily characterize the attack chain. However, when the subject/object of system events are different, events of the same

type may have different semantics. For example, we consider the file reading event, reading a downloaded file (R1) represents an untrusted data flow, while reading a text document in a user’s key directory (R2) represents a high-value data flow. The semantics of these two file reading events are not the same. In the ATT&CK model, R1 is mostly the behavioral characteristic in the Initial Access stage of the APT attack, which may lead to malicious code execution, while R2 may be the Data Exfiltration stage where the attacker is stealing user information, which may eventually lead to user information leakage. In order to solve the problems above, we define a series of entity labels, as well as transfer rules for labels between entities to visually describe the APT attack. In summary, **APTSHIELD gathers the semantics of the complete attack chain into target entity to realize the abstraction of large-scale attack features with constant-level memory overhead.**

4.5.2 Label Definition

(1) Process Labels. We divides process labels into two categories: status labels and behavior labels, as shown in Table 2. The status labels represent the label that semantics in the process will be passed to the child process with the FORK event. The behavior labels indicate what the process has done, and are used to accurately locate the threat. For example, if a process has network connection, we label it with *PS1*, for we can not trust the data from network. If a process executes a command without being allowed, we label this process with *PB4*. A detailed description of the process label is in Appendix .1.

Table 2. Process label definition; including status labels and behavior labels.

Labels	Description	Category
<i>PS1</i>	Process has network connection	Status
<i>PS2</i>	Process accesses to data in high-value data flow nodes	
<i>PS3</i>	Process contains network data	
<i>PS4</i>	Process loads or reads the uploaded file from network	
<i>PS5</i>	Process interacts with non-existent files	
<i>PS6</i>	Process reads the /etc/passwd file	
<i>PS7</i>	Process reads the .bash_history file	
<i>PB1</i>	Process executes a file from network	Behavior
<i>PB2</i>	Process executes a sensitive file	
<i>PB3</i>	Process executes a sensitive command	
<i>PB4</i>	Process has executed commands (for Webshell only)	
<i>PB5</i>	Process executes commands which can let others get the shell	
<i>PB6</i>	Process modifies the /etc/crontab file	
<i>PB7</i>	Process modifies the /etc/sudoers file	
<i>PB8</i>	Process reads high-value information	

(2) File Labels. We divide file labels into two categories: untrusted labels and high-value labels, as shown in Table 3. The untrusted files refer to the file containing untrusted data from the network, and the high-value files refer to the

file containing sensitive data. For example, if a file contains data from the network, we label it with *FU2*, for the untrusted data may cause an attack later (code execution). A detailed description of the file label is in Appendix .2.

Some processes with specific labels may directly or indirectly cause potential damage to system security. We define them as potential harmful functionalities (PHF). In this paper, we define *PS2, 3, 5* and *PB1, 2, 5* as PHF. Administrators are able to add more labels to the PHF list to implement strict control.

Table 3. File label definition; including untrusted labels and high-value labels.

Labels	Description	Category
<i>FU1</i>	File is uploaded	Untrusted
<i>FU2</i>	File contains data from the network	
<i>FU3</i>	File does not exist	
<i>FU4</i>	File is written by the Webshell attack	
<i>FU5</i>	File is written by the RAT attack	
<i>FU6</i>	File is written by the Living-off-the-land attack	
<i>FH1</i>	File that can control scheduled tasks such as /etc/crontab	High-value
<i>FH2</i>	File that can control permissions such as /etc/sudoers	
<i>FH3</i>	File that holds sensitive information such as /etc/passwd	
<i>FH4</i>	File that saves historical commands such as .bash_history	
<i>FH5</i>	File is written by process that have read sensitive information	

4.5.3 Event Selection

After defining the labels of system entities (i.e., processes and files), it is necessary to clarify effective events between entities in order to facilitate the transfer and aggregation of contextual information. Inspired by [27, 29], the events involved are a part of the events that exist in the Linux kernel data collected by auditd, including process events such as fork, execute, LoadELF, file operation, entity attribute modification, network operation, etc.. A list of the events we used is in Appendix .3.

4.5.4 Transfer Rules

To gather the semantics of the complete APT attack chain into target entity, we design a label transfer rule based on the ATT&CK model, as shown in Table 4. We divide APT attacks into five major stages, which are initial access, untrusted execution, lateral movement, suspicious behavior, and data exfiltration. Furthermore, the stage of suspicious behavior contains persistent stronghold, privilege escalation, credential access, and information collection. In Table 4, label 1 and label 2 represent semantic labels carried by system entities, and the direction represents the flow direction of system events, which is also the label transmission direction. For example, if a process with label *PS1* writes a file, we label the file with *FU2*, when a file with label *FU2* is accessed by a process, we label the process

with *PS3*. A detailed description of the transfer rules is in Appendix .4. If a certain process contains some labels in above five stages, the process has the possibility to represent an ongoing APT attack, the APT judge rules will be described in Section 4.5.5.

From the ATT&CK model and TTP, it is known that each stage of the APT attack represents a tactic, and a tactic (i.e., initial access, untrusted execution, etc.) contains different techniques. Analysts can expand Table 4 according to different techniques under specific tactics.

4.5.5 Judgement Rules

The attack judgment rules used in this paper are listed in Table 5. The last column specifies the prerequisites for the transfer rule to match. The prerequisites can specify conditions on the parameters of the alarm being matched. The administrator is able to judge whether an attack has occurred based on the different labels contained in the process. In other words, the labels of different entities on the attack chain will eventually come to a process through transfer and aggregation, when labels contained in the process match target judgment rules, the corresponding alerts (threats or APTs) will be reported to the administrator. For example, a process that contains the Webshell/RAT/Living-off-the-land attack labels, as well as the label *PB8* may be the alert point of APT attack. A detailed description of the judgment rules is in Appendix .5.

5. Design Experiments

In our evaluation, we first describe the experimental setup. Then we introduce the effectiveness of data compaction, the accuracy and effectiveness of APT detection, and the overhead of the system in turn.

5.1. Experimental Setup

Our datasets consist of two parts, one is collected from our laboratory, and the other is from Darpa Engagement. Table 6 summarizes the property of our dataset.

5.1.1 The Dataset from Laboratory

For the data from our laboratory, two participants (red team) in our laboratory were responsible for instrumenting OS and carrying out attack campaigns (on three hosts with Ubuntu 16.04), while the other two participants (blue team) performed data collection, data compaction and attack detection in a real-time manner. The benign background activities were also being carried out on the machines used for experimentation, such as web browsing, chatting, and document editing. In order to get an adversarial engagement, the blue team had no prior knowledge of the attacks planned by the red team. As shown in Table 6, L-1, L-2, and L-3

were the APT datasets generated by simulating real scenes in our laboratory, and the collection of each dataset lasted for 24 hours. Each dataset contained normal data and attack data. The attacks contained in above three datasets (L-1, L-2, and L-3) are: Webshell attack from the backdoored Apache, Remote Access Trojan from the phishing website, and Living-off-the-land attack from vulnerable service. A detailed description of attacks contained in L-1, L-2, and L-3 are in Appendix .6.

5.1.2 The Dataset from Darpa Engagement

For the data from Darpa Engagement, two hosts with Ubuntu 14.04 were deployed in advance for the data collection (blue team) and the APT attack (red team). Note that Blue team had no prior knowledge about the attack prepared by the red team. Similarity, when the red team was attacking the host, other activities of normal programs on the host were also proceeding simultaneously. Benign activities contained browsing websites, downloading and executing binary files, reading/writing emails and documents, etc.. Overall, more than 99% of the events in the dataset had nothing to do with attacks. As shown in Table 6, E-1 and E-2 are the APT datasets generated by Darpa Engagement, and the collection of each dataset lasted for 168 hours. Each dataset contained normal data and attack data. The attacks contained in above two datasets (E-1 and E-2) are: Information gather and exfiltration, Malicious file download and execute, and In-memory attack with firefox. A detailed description of attacks contained in E-1 and E-2 are in Appendix .7.

5.2. System Performance

Here, we will introduce the effectiveness of data compaction, the accuracy and effectiveness of APT detection, and the system overhead.

5.2.1 Performance of Data Compaction

(1) Effectiveness of Redundant Semantics Skipping

In real-time APT detection, the strategy of redundant semantics skipping is adopted to reduce events with repetitive semantics. It can be carried out in real-time streaming data to ensure the performance of the detection system (i.e., reduce memory overhead and running time). Table 7 gives results of redundant semantics skipping.

For the datasets (L-1, L-2, and L-3) from our laboratory, each dataset contains about 1000K events, and about 350K pieces of events can be skipped on average through the redundant semantic skipping strategy. In the real-time APT detection, APTSHIELD can save 1000ms-2500ms for every 1000K samples. For the datasets (E-1 and E-2) from Darpa Engagement, the datasets contain about 28000K-34000K events, and about 6000K pieces of events

Table 4. Transfer rules. Label 1 and label 2 represent semantic labels carried by system entities, and the direction represents the flow direction of system events. D stands for front to back, R stands for back to front.

Stage	Label1	Event	Label2	Direction	Description	
Initial Access	PS1	Write	FU2	D	A network-connected process wrote a file	
	PS3	Read	FU2	R	The process reads a file containing network data	
	PS3	Write	FU2	D	The process, which has accessed network data, writes these data to files	
	PS4	Read/Mmap	FU1	R	The file uploaded by the user is loaded or read by processes	
	PS5	Read/LoadElf/Mmap	FU3	R	The process interacts with non-existent files	
Untrusted Execution	PB1	Execute	FU2	R	The file, which contains the codes from the network, is executed	
	PB1	LoadElf	FU2	R	The file, which contains the codes from the network, is loaded	
	PB1	Write	FU2	D	The process, which has executed the codes from the network, writes files	
Lateral Movement	PB4	Write	FU4	D	The process, which has caused the Webshell attack, writes files	
	PB4	Read/Mmap	FU4	R	The process reads or loads files written by Webshell	
	PB1	Write	FU5	D	The process, which has caused the RAT attack, writes files	
	PB1	Read	FU5	R	The process reads files written by RAT	
	PS5	Write	FU6	D	The process, which has caused the Living-off-the-land attack, writes files	
Suspicious Behavior	PS5	Read	FU6	R	The process reads files written by Living-off-the-land	
	Persistent Stronghold	PB6	Write	FH1	R	The process writes files that can control scheduled tasks such as /etc/crontab
	Privilege Escalation	PB7	Write	FH2	R	The process writes files that can control permissions such as /etc/sudoers
	Credential Access	PS6	Read	FH3	R	The process reads files that hold sensitive information such as /etc/passwd
Data Exfiltration	Information Collection	PS7	Read	FH4	R	The process reads files that save historical commands such as .bash_history
	PB6 - 7 PS6 - 7	Write	FH5	D	The process writes high-value data to files	
	PB8	Read	FH5	R	The process reads files that contain high-value data	

Table 5. Judgement rules for different attacks (threats and APTs). The prerequisites can specify conditions on the parameters of the alarm being matched.

Alert	Condition	Prerequisites
Download&Execution	PB1	$init(p) = match(p.name, "/(scp wget)/") \rightarrow p.labels.add("PS1")$ $propRd(p, f) = if(p.labels.contains "PS1" and e == Write) \rightarrow f.labels.add("FU2")$ $propRd(f, p) = if(f.labels.contains "FU2" and e == Execute) \rightarrow p.labels.add("PB1")$ $alert(p) = if(p.labels.contains "PB1") \rightarrow Alert(Download&Execution)$
Webshell	PS4&PB4	$init(f) = match(f.name, "/(upload)/") \rightarrow p.labels.add("FU1")$ $propRd(f, p) = if(f.labels.contains "FU1" and e == ReaderMmap) \rightarrow p.labels.add("PS4")$ $inti(p) = match(p.name, "*") \rightarrow p.labels.add("PB4")$ $alert(p) = if(p.labels.contains "PS4&PB4") \rightarrow Alert(Webshell)$
RAT	PB1&PB5	$init(p) = match(p.name, "/(sh bash)/") \rightarrow p.labels.add("PB5")$ $inti(p) = if(p.cause.alert(Download&Execution)) \rightarrow p.labels.add("FU2")$ $alert(p) = if(p.labels.contains "PB1&PB5") \rightarrow Alert(RAT)$
Living-off-the-land	PS5&PB5	$init(f) = match(f.name, "/(null)/") \rightarrow f.labels.add("FU3")$ $propRd(f, p) = if(f.labels.contains "FU3" and e == LoadorRead) \rightarrow p.labels.add("PS5")$ $inti(p) = match(p.name, "/(sh bash)/") \rightarrow p.labels.add("PB5")$ $alert(p) = if(p.labels.contains "PS5&PB5") \rightarrow Alert(Living-off-the-land)$
Suspicious Behavior	PB6 PB7 PS6 PS7	$init(f) = match(f.name, "/etc/crontab") \rightarrow f.labels.add("FH1")$ $propRd(f, p) = if(f.labels.contains "FH1" and e == Write) \rightarrow p.labels.add("PB6")$ $init(f) = match(f.name, "/etc/sudoers") \rightarrow f.labels.add("FH2")$ $propRd(f, p) = if(f.labels.contains "FH2" and e == Write) \rightarrow p.labels.add("PB7")$ $init(f) = match(f.name, "/etc/passwd") \rightarrow f.labels.add("FH3")$ $propRd(f, p) = if(f.labels.contains "FH2" and e == Read) \rightarrow p.labels.add("PS6")$ $init(f) = match(f.name, ".bash_history") \rightarrow f.labels.add("FH4")$ $propRd(f, p) = if(f.labels.contains "FH2" and e == Read) \rightarrow p.labels.add("PS7")$ $alert(p) = if(p.labels.contains "PB6 PB7 PS6 PS7") \rightarrow Alert(SuspiciousBehavior)$
Data Exfiltration	PB8	$inti(p) = if(p.cause.alert(SuspiciousBehavior)) \rightarrow p.labels.add("PB6 PB7 PS6 PS7")$ $propRd(p, f) = if(p.labels.contains "PB6 PB7 PS6 PS7" and e == Write) \rightarrow f.labels.add("FH5")$ $propRd(f, p) = if(f.labels.contains "FH5" and e == Read) \rightarrow p.labels.add("PB8")$ $alert(p) = if(p.labels.contains "PB8") \rightarrow Alert(DataExfiltration)$
APT	((PS4&PB4) (PB1&PB5) (PS5&PB5))&PB8	$alert(p) = if(p.cause.alert(Webshell RAT Living-off-the-land) and p.labels.contains "PB8") \rightarrow Alert(APT)$

are able to be skipped on average through the redundant semantic skipping strategy. Although the proportion of skipped events has decreased compared with the data from laboratory (down from 35% to 20%), the time saved per 1000K events has increased significantly, reaching 7800ms-11000ms.

We find that APTSHIELD can save more time on the Darpa Engagement dataset, the reason is that the process of file reading and writing events takes more time, and there are a large number of such events on the Darpa Engagement dataset.

(2) Effectiveness of Non-viable Entity Pruning

Table 6. Details of our Datasets. L stands for data from laboratory, and E stands for data from Darpa Engagement.

Dataset	Duration (hours)	Platform	Source	Attack Description
L-1	24	Ubuntu 16.04 (64 bit)	Laboratory	Webshell attack from the backdoored Apache
L-2	24	Ubuntu 16.04 (64 bit)	Laboratory	Remote Access Trojan from the phishing website
L-3	24	Ubuntu 16.04 (64 bit)	Laboratory	Living-off-the-land attack from vulnerable service
E-1	168	Ubuntu 14.04 (64 bit)	Darpa Engagement	Information gather and exfiltration & Malicious file download and execute
E-2	168	Ubuntu 14.04 (64 bit)	Darpa Engagement	In-memory attack with firefox

Table 7. Results of Redundant Semantics Skipping.

Dataset	Sum of Events	Skip of Events	Saved Time (ms)/1000K
L-1	1142480	381885	1716
L-2	1070143	357307	1038
L-3	1006603	335721	2514
E-1	27978043	6019924	7790
E-2	33743425	6016741	11050

As the running time increases, incremental entities will be saved. However, some non-viable entities will neither pose a threat nor affect the forensic analysis. Therefore, we can delete these non-viable entities based on EXIT events and PHF labels to reduce memory consumption. Results of non-viable entities pruning are shown in Table 8. The results on different datasets show that with the growth of time and the increase of entities, the proportion of pruned processes is about 3% to 8%. Moreover, some inactive files will be moved to the disk (inactivity is defined as no change for more than 5 minutes) for reducing overhead. We can find from Table 8 that the pruned files mainly depend on the degree of interaction between files and processes in the dataset. The reason file pruning ratio of L-3 and E-1 is relatively high is because these two datasets contain a large number of inactive temporary files. These temporary files are all moved to disk. In addition, we also evaluate the CPU and memory overhead of APTSHIELD with/without pruning, which will be discussed in detail in Section 5.3.

Table 8. Results of Non-viable Entities Pruning. Sum stands for the total number of entities (processes or files). Remain stands for the number of entities (processes or files) remaining after being pruned.

Dataset	Prune for Process			Prune for File		
	Sum	Remain	Pruning Ratio (%)	Sum	Remain	Pruning Ratio (%)
L-1	125	121	3.2	144	138	4.2
L-2	91	84	7.7	255	247	3.1
L-3	127	119	6.3	138	53	61.6
E-1	16926	16226	4.1	1360	114	91.6
E-2	21673	20829	3.9	1232	1209	1.9

Table 9. Comparison with related work. L represents low, M represents medium, and H represents high.

Method	Real-time	Efficiency	Overhead
APTSHIELD	H	H	L
LCD (2016) [47]	H	L	L
FD (2018) [25]	L	M	H
GS (2021) [59]	M	H	H

(3) Comparison with Related Work

In APTSHIELD, we use a data compaction method that maintains global semantics to delete the equivalent information flow to the same target entity with low overhead by using the semantic attributes of the source entity. In addition, our APT detection and data compaction can be performed at the same time to ensure real-time performance. As shown in Table 9, we compared APTSHIELD to other existing data compaction methods from the following three aspects: real-time performance, efficiency, and overhead.

LCD [47] considers the event compaction of the single node in a real-time manner, but it does not link the contextual semantics, which could cause difficulty in forensics analysis. Moreover, LCD cannot handle network events. The input of FD [25] is the whole dependence graph, which can only start after data collection is completed (i.e., FD uses offline cached data). Also, FD cannot deal with events of subprocess/thread. GS [59] is able to establish global dependencies as FD does, but the compaction system of GS does not have a discarding mechanism (e.g., the processing of lengthy list fields, killed processes, and unused files). As time increases, the real-time performance and efficiency of the system will decrease. Our compaction method is able to improve the efficiency of real-time APT detection and reduce data storage overhead for forensic analysis by adopting redundant semantics skipping and non-viable entity pruning. All in all, the requirements of real-time performance, efficiency, and overhead can be jointly addressed by APTSHIELD. It outperforms all the above related methods.

5.2.2 Accuracy and Effectiveness of APT Detection

In this section, we will explain how the APT detection framework can effectively detect attacks in Table 6.

L-1: Webshell attack from the backdoored Apache.

The APT attack chain in dataset L-1 is shown in Figure 3. The process of the APT detection framework is as follows:

(1) The folder directory `/var/www/html/uploads/` in the web service is to save uploaded files. When the Webshell file `shell.php` has been uploaded to this directory by the Apache vulnerability, we label this file with `FU1` according to the file label definition (time point 1-3).

(2) If the file `shell.php` is accessed or loaded by a process, the process will be labeled with `PS4` by the transfer rule (`PS4, READ/MMAP, FU1, R`) (time point 4). Note that the Apache in Figure 3 is a subprocess, which will not cause label taint (i.e., other behaviors of apache will not be labeled).

(3) When a process with `PS4` executes any commands, the process will be labeled with `PB4` (time point 5-6). In order to track the spread of attacks after the invasion, APTSHIELD will label files accessed by the process (with `PB4`) as `FU4` according to the transfer rule (`PB4, WRITE, FU4, D`), which indicates that these files may be threatened by the subsequent proliferation of Webshell. When the file is accessed by other processes, the semantics will be passed back to the process (labeled with `PB4`) by the transfer rule (`PB4, READ/MMAP, FU4, R`), which makes subsequent infiltrations of Webshell fully recorded and detected.

(4) The process of Webshell (`PB4`) writes a controllable script (`/tmp/cleanup.sh`), (`/tmp/cleanup.sh`) will be labeled with `FU4` by the transfer rule (`PB4, WRITE, FU4, D`) (time point 9 and 16).

(5) The script (`FU4`) is loaded and executed by a privileged process (`/bin/sh`), (`/bin/sh`) will be labeled with `PB4` by the transfer rule (`PB4, READ/MMAP, FU4, R`), (`/bin/sh`) executes `nc` to create command&control between the user and the attacker (time point 17-23).

(6) The process `sh` (`PB4`) continues to access sensitive files, such as `/etc/crontab`(`FH1`), `/etc/sudoers`(`FH2`), `/etc/passwd`(`FH3`), and `.bash_history`(`FH4`), it will be labeled with the corresponding label (i.e., `PB6`, `PB7`, `PS6`, and `PS7`) indicating that it has read or written these sensitive files (time point 26-35).

(7) When a process `sh` has read or written sensitive files that wants to pass high-value information out, the file `secret` written by it will be labeled with `FH5` (time point 36-38), indicating that the file `secret` may contain high-value information (`PS6-7/PB6-7, WRITE, FH5, D`).

(8) If another process `cat` subsequently accesses the file with the `FH5` label, it may be labeled with `PB8` due to the threat of data exfiltration (`PB8, READ, FH5, R`) (time

point 39-40). Finally, the process that satisfies this series of attack chains may generate the threat of APT attacks (including Initial Access, Lateral Movement, Suspicious Behavior, and Data Exfiltration). Therefore, it is reported that there may be APT which uses Webshell as the entrance to the process that contains labels of Webshell and `PB8` at the same time.

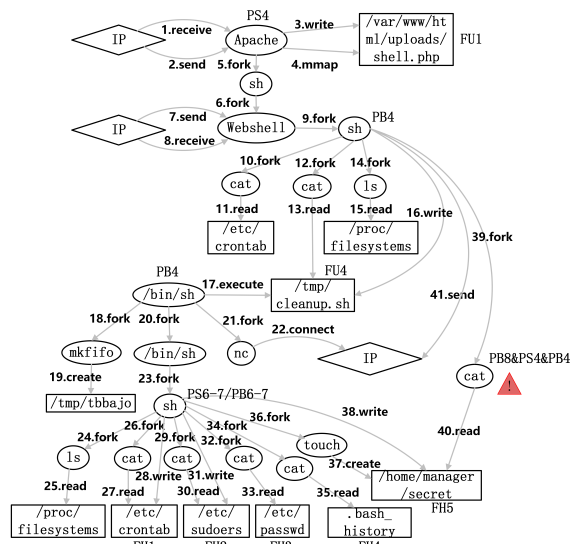


Figure 3. APT attack chain that uses Webshell as the entrance.

L-2: Remote Access Trojan from the phishing website.

For the APT in L-2, the detection flow of APTSHIELD is as follows:

(1) Through a phishing link, the target host downloads the Trojan file (`vpn.elf`). Since the file is transmitted over the Internet, it will be labeled with `FU2`.

(2) After a period of time, the file (`vpn.elf`) is executed, and the target process will be labeled with `PS1` by the transfer rule (`PS1, EXECUTE, FU2, R`), indicating that the process executes files originating from the Internet may cause a threat.

(3) Subsequent attacks and proliferation are the same as they are in the Webshell experiment. Finally, it is reported that there may be APT which uses RAT as the entrance to the process that contains labels of RAT and `PB8` at the same time.

L-3: Living-off-the-land attack from vulnerable service.

For the APT in L-3, the detection flow of APTSHIELD is as follows:

(1) Through an exploit script, the threatening process will interact with a file named "`null`". Therefore, the non-existent file will be labeled with `FU3`.

(2) When a process interacts with the file, the process will be labeled with `PS5` by the transfer rule (`PS5, READ/MMAP/LOADELF, FU3, R`). When

the process containing the label *PS5* continues to execute the *sh* or *bash* command, it will be labeled with *PB5*, which means that the process has got the shell and carried out a in-memory attack.

(3) Subsequent attacks and proliferations are the same as they are in the Webshell experiment. Finally, it is reported that there may be APT that uses Living-off-the-land as the entrance for the process that contains labels of Living-off-the-land and *PB8* at the same time.

E-1: Information gather and exfiltration. The detection flow of APTSHIELD is as follows:

(1) The file (*/etc/passwd*) containing sensitive information are labeled with *FH3*.

(2) When a process reads this file, the process will be labeled with *PS6* by the transfer rule (*PS6, READ, FH3, R*).

(3) The process writes the high-value data on */dev/pts/1*, */dev/pts/1* will be labeled with *FH5* by the transfer rule (*PS6, WRITE, FH5, D*). By this way, the user may have exfiltrated this information.

E-1: Malicious file download and execute. The attack chain is shown in Figure 4. The detection flow of APTSHIELD is as follows:

(1) A file named *ccleaner* is downloaded to the local through *scp*, this file will be labeled with *FU2*. Because the process that uses the *scp* command is likely to include an internet connection *PS1*, the file will be labeled with *FU2* by the transfer rule (*PS1, WRITE, FU2, D*).

(2) After the file is executed by the process, the process will be labeled with *PB1* by the transfer rule (*PB1, EXECUTE, FU2, R*). Therefore, the process containing the *PB1* label may pose a threat (download and execution).

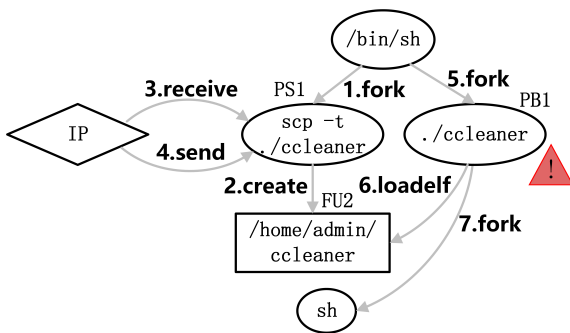


Figure 4. APT attack chain that uses *ccleaner* as the entrance.

E-2: In-memory attack with firefox. The attack chain is shown in Figure 5. The detection flow of APTSHIELD is as follows:

(1) Use Firefox’s write-executable memory space to cause in-memory attacks through *sshd* process. The *sshd* process interacts with the *"/(null)"* file, and

the process is labeled with *PS5* by the transfer rule (*PS5, READ/MMAP/LOADELF, FU3, R*).

(2) Later, the process with *PS5* gets the shell, it will be labeled with (*PB5*). Processes that contain the *PB5* and *PS5* tags may cause the threat of in-memory attacks.

(3) Finally, a file named *hc* was uploaded and executed through the *sshd* in-memory attack. The detection flow is the same as the *ccleaner* attack.

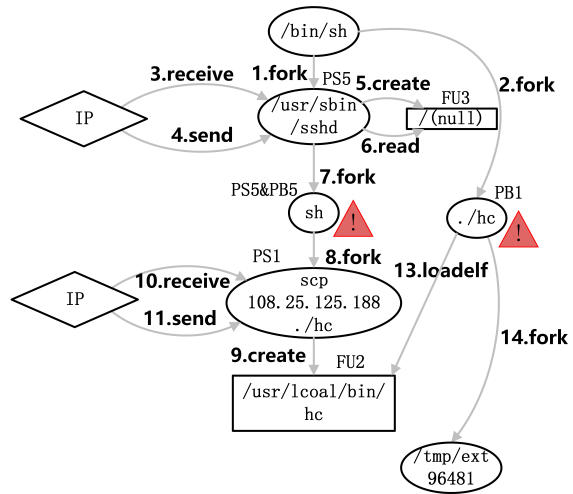


Figure 5. APT attack chain that uses in-memory attack as the entrance.

5.3. Overhead of APTSHIELD

5.3.1 CPU and Memory Usage

The CPU and memory usage of APTSHIELD for five different datasets are shown in Table 10. We use performance monitor to record the average overhead of the system while it is running. Here, the CPU usage means the usage rate on a single-core CPU (thread). In the absence of pruning, the average single-core CPU usage of APTSHIELD is 5.3%, 5.8%, 4.7%, 9.7%, and 10.1% on L-1, L-2, L-3, E-1, and E-2, respectively, while the average memory usage is 90.1 MB, 99.7 MB, 85.6 MB, 117.7 MB, and 130.3 MB on L-1, L-2, L-3, E-1, and E-2, respectively. By deploying non-viable entity pruning (to further reduce the system overhead), the average single-core CPU usage is reduced by 1.1%, 0.7%, 0.4%, 2.0%, and 1.6% on L-1, L-2, L-3, E-1, and E-2, respectively, while the average memory usage is reduced by 14.8 MB, 22.6 MB, 14.9 MB, 22.2 MB, and 19.7 MB on L-1, L-2, L-3, E-1, and E-2, respectively. All in all, in the case of real-time detection, APTSHIELD can maintain low CPU usage and satisfactory memory overhead.

Table 10. The average overhead of APTSHIELD. By deploying non-viable entity pruning, the overheads of CPU and memory are reduced.

Dataset	Without Pruning		With Pruning	
	CPU (%)	Memory (MB)	CPU (%)	Memory (MB)
L-1	5.3	90.1	4.2%	75.3
L-2	5.8	99.7	5.1%	77.1
L-3	4.7	85.6	4.3%	70.7
E-1	9.7	117.7	7.7%	95.5
E-2	10.1	130.3	8.5%	110.6

5.3.2 Real-Time Performance

We test the real-time performance of the APTSHIELD by comparing data consumption rate and data generation rate. In the experiment, the personnel in our laboratory performed normal operations on Linux host for 600 seconds, making the collected content as similar as possible to five datasets (L-1, L-2, L-3, E-1, and E-2). Results are shown in Figure 6.




Figure 6. Real-time Performance of APTSHIELD. Comparison of data consumption rate and data generation rate.

In Figure 6, the line on the right shows that about 230K pieces of data (contain entities and events) are generated and transmitted in real-time within ten minutes. The smooth phase of the line in Figure 6 stands for the idle state of the host, while the steep phase of the line stands for the busy state of the host, and the transmission rate is about 400 per second. The other line represents the consumption rate of APTSHIELD. About 213K pieces of data can be processed within 5 seconds, and the average processing rate is about 42K per second. The consumption rate is far greater than the generation rate, which shows the real-time of APTSHIELD is feasible.

5.3.3 Comparison with Related Work

We compare APTSHIELD with state-of-the-art studies to show the advantage of our work.

(1) **Conan** [29]. Conan only considers the data compaction in the storage stage, while APTSHIELD can perform compaction and detection at the same time. In addition,

APTSHIELD also prunes invalid/unused entities to reduce overhead. In our experiment, the average memory usage of Conan on the dataset from laboratory (L-1, L-2, and L-3) is about 90 MB. For the dataset from Darpa Engagement, the average memory usage of Conan is about 140 MB. Compared with APTSHIELD, the memory overhead of Conan is increased by about 30% to 50%, and the growth rate of CPU usage reaches to 10% to 20%. Furthermore, APTSHIELD has a better semantic interpretation and good detection efficiency for in-memory attacks.

(2) **Sleuth** [27] & **Holmes** [28]. According to the description in Section 2, Sleuth will read all the data into the memory at one time during detection, which cannot ensure real-time performance. The memory Sleuth used is related to the size of the dataset. In our experiment, the memory usage of Sleuth on L-1, L-2 and L-3 ranges from 80 MB to 160 MB, while on E-1 and E-2, it ranges from 300 MB to 400 MB. For Holmes, authors in [28] showed that a nearly linear growth in memory consumption. In our experiment, the memory usage of Holmes on L-1, L-2 and L-3 ranges from 90 MB to 150 MB, while on E-1 and E-2, it ranges from 500 MB to 1 GB. Compared with APTSHIELD, the memory overhead of Sleuth and Holmes is increased by at least 60% and 20%, respectively. APTSHIELD has a good real-time performance and can quickly respond to threats that occur in the system.

6. Conclusion & Discussion

In this paper, we design a stable, efficient, and real-time APT detection system, called APTSHIELD. Unlike previous studies with the problems of the weak data source integrity, large data processing overhead and poor real-time performance, APTSHIELD can collect reliable, stable and semantically rich data sources via audit, reduce the overhead of the detection system in real-time via redundant semantics skipping and non-viable node pruning, and carry out real-time APT attack response and alarm through the transfer and aggregation of labels based on ATT&CK model. Experimental results on both laboratory and Darpa Engagement show that APTSHIELD can effectively detect web vulnerability attacks, file-less attacks and remote access trojan attacks, at the same time has a low false positive rate, which adds far more value than the existing frontier work. And we will discuss the limitations and improvements of our work.

6.1. False Positive

Although APTSHIELD did not show any false positives in the experiment of this article, we still need to discuss the following scenarios:

Normal Meta behavior. Since an APT attack is composed of multiple different behaviors, each individual behavior (we call it meta behavior) does not represent the oc-

currence of an attack. For example, *READ* events to high-value files such as */etc/passwd* may occur during the APT attack stage (when the user is authenticated, */etc/passwd* will be accessed), but related processes should not be identified as threats. In order to prevent single-point false positives, APTSHIELD adopts the method of contextual semantic transfer to detect the APT attack chain. Analysts can set stringent judgment rules to avoid false positives.

Taint explosion. Recall that through the transfer and aggregation of labels, we are able to correlate the context information of the attack. However, this method is prone to cause taint explosion (e.g., polluting the daemon process of the system at the entry point, causing all entities interacting with the process to be polluted). A lot of false positives will occur due to taint explosion. In order to solve this problem, we can formulate stricter labeling and transfer rules to ensure that the daemon process is not easily marked as suspicious when it has not been invaded or modified.

6.2. False Negative

Reducing the false negative rate is a major challenge for APT detection, because APT attacks are completely unknown. APTSHIELD uses different tactical stages in the ATT&CK model to describe a complete APT attack. In APT detection, APTSHIELD can reflect the nature of the attack by tracking untrusted data flow, untrusted control flow, and high-value data flow. However, for some alternative attacks, such as open source software supply chain attacks (i.e., to inject malicious code into open source software packages in order to compromise dependent systems further down the chain) [62]. Some well-known open source softwares may be whitelisted by system administrators, allowing them to be released at the system entry point. To solve this problem, we can combine with code dependency analysis at the entry point to label unsafe factors and reduce false negatives.

References

- [1] What Is Stuxnet. <https://www.mcafee.com/enterprise/en-hk/security-awareness/ransomware/what-is-stuxnet.html>. 1
- [2] BlackEnergy APT Attacks in Ukraine. <https://www.kaspersky.com/resource-center/threats/blackenergy>. 1
- [3] Missed Alarms and 40 Million Stolen Credit Card Numbers: How Target Blew It. <https://www.bloomberg.com/news/articles/2014-03-13/target-missed-warnings-in-epic-hack-of-credit-card-data>. 1
- [4] APT trends report Q3 2020. <https://securelist.com/apt-trends-report-q3-2020/99204/>. 1
- [5] APT Groups Target Firms Working on COVID-19 Vaccines. <https://www.bankinfosecurity.com/microsoft-warning-a-15363>. 1
- [6] Alexander D Bolton and Christine M Anderson-Cook. Apt malware static trace analysis through bigrams and graph edit distance. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 10(3):182–193, 2017. 1, 4
- [7] Giuseppe Laurenza, Leonardo Aniello, Riccardo Lazzeretti, and Roberto Baldoni. Malware triage based on static features and public apt reports. In *International Conference on Cyber Security Cryptography and Machine Learning*, pages 288–305. Springer, 2017. 1, 4
- [8] Donglan Liu, Hao Zhang, Hao Yu, Xin Liu, Yong Zhao, and Guodong Lv. Research and application of apt attack defense and detection technology based on big data technology. In *2019 IEEE 9th International Conference on Electronics Information and Emergency Communication (ICEIEC)*, pages 1–4. IEEE, 2019. 1, 4
- [9] Ishai Rosenberg, Guillaume Sicard, and Eli Omid David. Deepapt: nation-state apt attribution using end-to-end deep neural networks. In *International Conference on Artificial Neural Networks*, pages 91–99. Springer, 2017. 1, 4
- [10] Guodong Zhao, Ke Xu, Lei Xu, and Bo Wu. Detecting apt malware infections based on malicious dns and traffic analysis. *IEEE access*, 3:1132–1142, 2015. 1, 4
- [11] Lu Huang, Jingfeng Xue, Weijie Han, Zixiao Kong, and Zequn Niu. Detection of malicious domains in apt via mining massive dns logs. In *International Conference on Machine Learning for Cyber Security*, pages 140–152. Springer, 2020. 1, 4
- [12] Natasha Arjumand Shoaib Mirza, Haider Abbas, Farukh Aslam Khan, and Jalal Al Muhtadi. Anticipating advanced persistent threat (apt) countermeasures using collaborative security mechanisms. In *2014 International Symposium on Biometrics and Security Technologies (ISBAST)*, pages 129–132. IEEE, 2014. 1, 4
- [13] Amin Kharaz, Sajjad Arshad, Collin Mulliner, William Robertson, and Engin Kirda. {UNVEIL}: A large-scale, automated approach to detecting ransomware. In *25th {USENIX} Security Symposium ({USENIX} Security 16)*, pages 757–772, 2016. 1, 4
- [14] What Is Endpoint Detection and Response. <https://www.mcafee.com/enterprise/zh-cn/security-awareness/endpoint/what-is-endpoint-detection-and-response.html>. 1
- [15] 2020 Gartner Market Guide for Network Detection and Response. <https://www.gigamon.com/resources/resource-library/analyst-industry-reports/2020-gartner-market-guide-for-network-detection-and-response.html>. 1
- [16] Transparent Computing Program. <https://github.com/darpa-i2o/Transparent-Computing>. 2
- [17] Trustwave Global Security Report. <https://www.trustwave.com/en-us/resources/library/documents/2020-trustwave-global-security-report/>. 2
- [18] Kyu Hyung Lee, Xiangyu Zhang, and Dongyan Xu. Loggc: garbage collecting audit log. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 1005–1016, 2013. 2, 3, 5

- [19] Kyu Hyung Lee, Xiangyu Zhang, and Dongyan Xu. High accuracy attack provenance via binary-based execution partition. In *NDSS*, 2013. 2, 4
- [20] Shiqing Ma, Xiangyu Zhang, and Dongyan Xu. Protracer: Towards practical provenance tracing by alternating between logging and tainting. In *NDSS*, 2016. 2, 4
- [21] Shiqing Ma, Juan Zhai, Fei Wang, Kyu Hyung Lee, Xiangyu Zhang, and Dongyan Xu. {MPI}: Multiple perspective attack investigation with semantic aware execution partitioning. In *26th {USENIX} Security Symposium ({USENIX} Security 17)*, pages 1111–1128, 2017. 2, 4
- [22] Yonghwi Kwon, Fei Wang, Weihang Wang, Kyu Hyung Lee, Wen-Chuan Lee, Shiqing Ma, Xiangyu Zhang, Dongyan Xu, Somesh Jha, Gabriela F Ciocarlie, et al. Mci: Modeling-based causality inference in audit logging for attack investigation. In *NDSS*, 2018. 2, 4
- [23] Yonghwi Kwon, Dohyeong Kim, William Nick Sumner, Kyungtae Kim, Brendan Saltaformaggio, Xiangyu Zhang, and Dongyan Xu. Ldx: Causality inference by lightweight dual execution. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 503–515, 2016. 2, 4
- [24] Wajih Ul Hassan, Lemay Aguse, Nuraini Aguse, Adam Bates, and Thomas Moyer. Towards scalable cluster auditing through grammatical inference over provenance graphs. In *Network and Distributed Systems Security Symposium*, 2018. 2
- [25] Md Nahid Hossain, Junao Wang, Ofir Weisse, R Sekar, Daniel Genkin, Boyuan He, Scott D Stoller, Gan Fang, Frank Piessens, Evan Downing, et al. Dependence-preserving data compaction for scalable forensic analysis. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*, pages 1723–1740, 2018. 2, 4, 5, 8, 13
- [26] Wajih Ul Hassan, Adam Bates, and Daniel Marino. Tactical provenance analysis for endpoint detection and response systems. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 1172–1189. IEEE, 2020. 2, 4
- [27] Md Nahid Hossain, Sadegh M Milajerdi, Junao Wang, Birhanu Eshete, Rigel Gjomemo, R Sekar, Scott Stoller, and VN Venkatakrishnan. {SLEUTH}: Real-time attack scenario reconstruction from {COTS} audit data. In *26th {USENIX} Security Symposium*, pages 487–504, 2017. 2, 5, 10, 16
- [28] Sadegh M Milajerdi, Rigel Gjomemo, Birhanu Eshete, Ramachandran Sekar, and VN Venkatakrishnan. Holmes: real-time apt detection through correlation of suspicious information flows. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1137–1152. IEEE, 2019. 2, 5, 16
- [29] Chunlin Xiong, Tiantian Zhu, Weihao Dong, Linqi Ruan, Runqing Yang, Yan Chen, Yueqiang Cheng, Shuai Cheng, and Xutong Chen. Conan: A practical real-time apt detection system with high accuracy and efficiency. *IEEE Transactions on Dependable and Secure Computing*, 2020. 2, 4, 5, 10, 16
- [30] MITRE ATT&CK. <https://attack.mitre.org/>. 2
- [31] Jiazhong Lu, Xiaosong Zhang, Wang Junfeng, and Ying Lingyun. Apt traffic detection based on time transform. In *2016 International Conference on Intelligent Transportation, Big Data & Smart City (ICITBS)*, pages 9–13. IEEE, 2016. 3
- [32] Branka Stojanović, Katharina Hofer-Schmitz, and Ulrike Kleb. Apt datasets and attack modeling for automated detection methods: A review. *Computers & Security*, 92:101734, 2020. 3
- [33] Yunfei Su, Mengjun Li, Chaojing Tang, and Rongjun Shen. A framework of apt detection based on dynamic analysis. In *2015 4th National Conference on Electrical, Electronics and Computer Engineering*, pages 1047–1053. Atlantis Press, 2015. 3
- [34] Yuan Wang, Yongjun Wang, Jing Liu, and Zhijian Huang. A network gene-based framework for detecting advanced persistent threats. In *2014 Ninth International Conference on P2P, Parallel, Grid, Cloud and Internet Computing*, pages 97–102. IEEE, 2014. 3
- [35] Jaeyeon Jung and Yu Zhu. Sensitive data tracking using dynamic taint analysis, Nov. 18 2014. US Patent 8,893,280. 3
- [36] Stephen Yang, Seo Jin Park, and John Ousterhout. Nanolog: A nanosecond scale logging system. In *2018 {USENIX} Annual Technical Conference*, pages 335–350, 2018. 3
- [37] Apache Log4j 2. <http://logging.apache.org/log4j/2.x/>. 3
- [38] Spdlog: A Super fast C++ logging library. <https://github.com/gabime/spdlog>. 3
- [39] glog: Google Logging Module. <https://github.com/google/glog>. 3
- [40] Boost C++ libraries. <https://github.com/google/glog>. 3
- [41] Auditd. <https://linux.die.net/man/8/auditd>. 3
- [42] Sysdig. <https://github.com/draios/sysdig>. 3
- [43] Ftrace. <https://web.ftrace.com>.
- [44] Ltng. <https://ltng.org>. 3
- [45] Auditbeat. <https://www.elastic.co/cn/beats/auditbeat>. 3
- [46] Yutao Tang, Ding Li, Zhichun Li, Mu Zhang, Kangkook Jee, Xusheng Xiao, Zhenyu Wu, Junghwan Rhee, Fengyuan Xu, and Qun Li. Nodemerge: Template based efficient data reduction for big-data causality analysis. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 1324–1337, 2018. 4
- [47] Zhang Xu, Zhenyu Wu, Zhichun Li, Kangkook Jee, Junghwan Rhee, Xusheng Xiao, Fengyuan Xu, Haining Wang, and Guofei Jiang. High fidelity data reduction for big data security dependency analyses. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 504–516, 2016. 4, 13
- [48] Zhenyuan Li, Qi Alfred Chen, Runqing Yang, Yan Chen, and Wei Ruan. Threat detection and investigation with system-level provenance graphs: a survey. *Computers & Security*, page 102282, 2021. 4
- [49] Xiaoyan Sun, Jun Dai, Peng Liu, Anoop Singhal, and John Yen. Using bayesian networks for probabilistic identification of zero-day attack paths. *IEEE Transactions on Information Forensics and Security*, 13(10):2506–2521, 2018. 4

- [50] Xueyuan Han, Thomas Pasquier, Adam Bates, James Mickens, and Margo Seltzer. Unicorn: Runtime provenance-based detector for advanced persistent threats. *arXiv preprint arXiv:2001.01525*, 2020. 4
- [51] Samuel T King and Peter M Chen. Backtracking intrusions. In *Proceedings of the nineteenth ACM symposium on Operating systems principles*, pages 223–236, 2003. 4, 5
- [52] Yushan Liu, Mu Zhang, Ding Li, Kangkook Jee, Zhichun Li, Zhenyu Wu, Junghwan Rhee, and Prateek Mittal. Towards a timely causality analysis for enterprise security. In *NDSS*, 2018. 4
- [53] Wajih Ul Hassan, Shengjian Guo, Ding Li, Zhengzhang Chen, Kangkook Jee, Zhichun Li, and Adam Bates. Nodoze: Combatting threat alert fatigue with automated provenance triage. In *Network and Distributed Systems Security Symposium*, 2019. 4
- [54] Shiqing Ma, Juan Zhai, Yonghwi Kwon, Kyu Hyung Lee, Xiangyu Zhang, Gabriela Ciocarlie, Ashish Gehani, Vinod Yegneswaran, Dongyan Xu, and Somesh Jha. Kernel-supported cost-effective audit logging for causality tracking. In *2018 {USENIX} Annual Technical Conference*, pages 241–254, 2018. 4
- [55] The Numbers Game: How Many Alerts are too Many to Handle? <https://www.fireeye.com/offers/rpt-idc-the-numbers-game.html>. 5
- [56] Mathieu Barre, Ashish Gehani, and Vinod Yegneswaran. Mining data provenance to detect advanced persistent threats. In *11th International Workshop on Theory and Practice of Provenance (TaPP 2019)*, 2019. 4
- [57] Ghita Berrada and James Cheney. Aggregating unsupervised provenance anomaly detectors. In *11th International Workshop on Theory and Practice of Provenance (TaPP 2019)*, 2019. 4
- [58] Sadegh M Milajerdi, Birhanu Eshete, Rigel Gjomemo, and VN Venkatakrisnan. Poirot: Aligning attack behavior with kernel audit records for cyber threat hunting. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 1795–1812, 2019. 5
- [59] Tiantian Zhu, Jiayu Wang, Linqi Ruan, Chunlin Xiong, Jinkai Yu, Yaosheng Li, Yan Chen, Mingqi Lv, and Tieming Chen. General, efficient, and real-time data compaction strategy for apt forensic analysis. *IEEE Transactions on Information Forensics and Security*, 2021. 4, 5, 8, 13
- [60] RSA loses SecureID tokens due to APT attack. <https://www.secplicity.org/2011/03/20/rsa-loses-secureid-tokens-due-to-apt-attack>. 6
- [61] UnixBench. <https://github.com/kdlucas/byte-unixbench>. 7
- [62] Marc Ohm, Henrik Plate, Arnold Sykosch, and Michael Meier. Backstabber’s knife collection: A review of open source software supply chain attacks. In *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*, pages 23–43. Springer, 2020. 17

.1. Explanation of Defined Label in Table 2

PS1: If a process has network connection, we label it with *PS1*, for we can't trust the data from network.

PS2: If a process has accessed high-value data flow nodes, we label it with *PS2*, for the process holds sensitive data.

PS3: If a process has data from the network, we label it with *PS3*, e.g., the process reads the file downloaded from network.

PS4: If a process has loaded or read the uploaded file from network, which can cause Webshell attack, we label it with *PS4* for the uploaded file can't be trusted.

PS5: If a process has interacted with non-existent files, which can cause Living-off-the-land attack, we label it with *PS5* for accessing to non-existent files is a feature of Living-off-the-land attacks.

PS6: If a process has read the file, which contains sensitive information such as */etc/passwd*, we label it *PS6* for it may cause leakage of sensitive information.

PS7: If a process has read the file, which saves historical commands such as *.bash_history*, we label it with *PS7* for the traces may be exploited.

PB1: If a process has executed a file from network, we label the process with *PB1*, for the file can't be trusted.

PB2: If a process has executed a sensitive file, we label the process with *PB2*, for the file may cause attack.

PB3: If a process has executed a sensitive command, we label the process with *PB3*, for the command may cause attack.

PB4: If a process has executed any command, we label the process with *PB4*, when the process is not allowed to execute commands (for Webshell only).

PB5: If a process has executed shell command, we label the process with *PB5*, for the process may reverse a shell to others.

PB6: If a process has modified the file that control scheduled tasks such as */etc/crontab*, we label the process with *PB6*, for it may cause persistent control.

PB7: If a process has modified the file that control permissions such as */etc/sudoers*, we label the process with *PB7*, for it may cause privilege escalation.

PB8: If a process has read high-value information, we label the process with *PB8*, for it may cause data exfiltration.

.2. Explanation of Defined Label in Table 3

FU1: If a file is uploaded, we label the file with *FU1*, for it can be untrustworthy.

FU2: If a file contains data from the network, we label the file with *FU2*, for the untrusted data may cause an attack later (code execution).

FU3: If a file does not exist, we label it with *FU3*. In a Living-off-the-land attack, accessing to a non-existent file

is marked as access to a (*null*) file.

FU4: If a file is written by the Webshell attack, we label the file with *FU4*, in order to capture the attack chain with webshell as the entry point.

FU5: If a file is written by the RAT attack, we label the file with *FU5*, in order to capture the attack chain with RAT as the entry point.

FU6: If a file is written by the Living-off-the-land attack, we label the file with *FU6*, in order to capture the attack chain with Living-off-the-land as the entry point.

FH1: If a file can control scheduled tasks such as */etc/crontab*, we label it with *FH1*.

FH2: If a file can control permissions such as */etc/sudoers*, we label it with *FH2*.

FH3: If a file holds sensitive information such as */etc/passwd*, we label it with *FH3*.

FH4: If a file saves historical commands such as *.bash_history*, we label it with *FH4*.

FH5: If a file is written by the process that have read sensitive information, we label the file with *FH5*, for it may cause data exfiltration.

.3. Event Used in APTSHIELD and Explanation

Table 11. Event definition used in APTSHIELD.

Events	Description
E0	File Read
E1	File Write
E2	Fork
E3	Execute
E4	LoadLibrary
E5	File Delete
E6	File Rename
E7	File Create
E8	File Property
E9	Exit
E10	LoadElf
E11	File Open
E12	File Close
E13	Fork with shared open file
E14	File Open with close-on-exec mark
E15	File Mmap

.4. Explanation of Defined Transfer Rules in Table 4

PS1-E1-FU2-D: If a process with label *PS1* writes a file, we label the file with *FU2*.

PS3-E0-FU2-R: If a file with label *FU2* is accessed by a process, we label the process with *PS3*.

PS3-E1-FU2-D: If a process with label *PS3* writes a file, we label the file with *FU2*.

PS4-E0/E15-FU1-R: If a file with label *FU1* is read or loaded by a process, we label the process with *PS4*.

PS5-E0/E10/E15-FU3-R: If a file with label *FU3* is interacted by a process, we label the process with *PS5*.

PB1-E3-FU2-R: If a file with label *FU2* is executed by a process, we label the process with *PB1*.

PB1-E10-FU2-R: If a file with label *FU2* is loaded by a process, we label the process with *PB1*.

PB1-E1-FU2-D: If a process with label *PB1* writes a file, we label the file with *FU2*.

PB4-E1-FU4-D: If a process with label *PB4* writes a file, we label the file with *FU4*.

PS4-E0/E15-FU4-R: If a file with label *FU4* is read or loaded by process, we label the process with *PS4*.

PB1-E1-FU5-D: If a process with label *PB1* writes a file, we label the file with *FU5*.

PB1-E0-FU5-R: If a file with label *FU5* is accessed by a process, we label the process with *PB1*.

PS5-E1-FU6-D: If a process with label *PS5* writes a file, we label the file with *FU6*.

PS5-E0-FU6-R: If a file with label *FU6* is accessed by a process, we label the process with *PS5*.

PB6-E1-FH1-R: If a file with *FH1* is written by a process, we label the process with *PB6*.

PB7-E1-FH2-R: If a file with *FH2* is written by a process, we label the process with *PB7*.

PS6-E0-FH3-R: If a file with *FH3* is accessed by a process, we label the process with *PS6*.

PS7-E0-FH4-R: If a file with *FH4* is accessed by a process, we label the process with *PS7*.

PB6-7|PS6-7-E1-FH5-D: If a process with any one of labels *PS1-4* writes a file, we label the file with *FH5*.

PB8-E0-FH5-R: If a file with label *FH5* is accessed by a process, we label the process with *PB8*.

.5. Explanation of Defined Judge Rules in Table 5

Download&Execution: If a process has a network connection and downloads files from the network, executing the file indicates that a threat of Download&Execution may have occurred. The label *PB1* represents that a process has the semantics of Download&Execution.

Webshell: If a process reads/loads the uploaded file and executes any commands without being allowed. It indicates that a Webshell attack may have occurred. The label *PB1* represents that a process has the semantics of Webshell attack.

RAT: If a process executes a file downloaded from the network and executes a shell command, which reverses a shell to others, it indicates that a RAT attack may have occurred. The label *PB1&PB5* represents that a process has the semantics of RAT attack.

Living-off-the-land: If a process interacts with a file that does not exist, and executes a shell command to reverse the shell to others, it indicates that a Living-off-the-land attack may have occurred. The label *PS5&PB5* represents that a process has the semantics of Living-off-the-land attack.

Suspicious Behavior: If a process reads or writes the file that controls permissions or holds sensitive information, it indicates that suspicious Behaviors (i.e., persistent stronghold, privilege escalation, credential access, information collection, etc.) may have occurred. The label *PB6|PB7|PS6|PS7* represents that a process has the semantics of suspicious Behaviors.

Data Exfiltration: If a process causes suspicious behavior and sends data out, it may cause the threat of data leakage. The label *PB8* represents that a process has the semantics of data exfiltration.

APT: Satisfy any one of Webshell attacks, RAT attacks, Living-off-the-land attacks, and satisfies data exfiltration at the same time. In the case of such an attack chain, APT attacks may exist. The label $((PS4&PB4)|(PB1&PB5)|(PS5&PB5))\&PB8$ represents that a process has the semantics of APTs.

.6. Details of Attacks from Laboratory

L-1: Webshell attack from the backdoored Apache. In this attack, port 80 is opened on the host to provide web services, and the website has file upload vulnerabilities while any file can be uploaded and accessed. The attacker successfully obtains the Webshell by uploading the Trojan horse file and accessing and connecting through Ant Sword. But the Webshell only has the www-data user permission. In the subsequent penetration, it is found that in the */etc/crontab* file, the administrator uses the script *cleanup.sh* to clean up the folder where uploaded files are stored once an hour. By viewing the script *cleanup.sh*, the attacker finds that anyone has the permission to modify it, so the attacker adds an attack command in *cleanup.sh* to get the reverse shell. Then the attacker starts to monitor and wait for the script to be executed by the crontab. Soon the attacker obtained the permission of the administrator. By using *id* command, the attacker finds that the administrator belongs to the root group and has all permissions required by the attacker.

Through the shell of the administrator, the attacker can modify */etc/crontab* and */etc/sudoers* to keep persistent and exploit benign users. Also, the attacker can check */etc/passwd* and *.bash_history* to obtain other user's credentials and collect their information. After that, the attacker writes the valuable information into a file and finally sends.

L-2: Remote Access Trojan from the phishing website. In this attack, the attacker uses a phishing website to put the Trojan file into the target host. After the monitoring is turned on and the Trojan file is executed, the per-

mission of benign user are obtained. Afterwards, same as the follow-up operation in the Webshell attack scenario, the timing script in the crontab is found, and the script is modified to obtain higher user's permissions. After that, the sensitive file is modified to achieve the purpose of persistence control and escalation. The sensitive file and user credentials are read, written into the file and finally sends.

L-3: Living-off-the-land attack from vulnerable service. In this attack, the attacker finds that a special service is running on port 29273 on the host, and the permission of users could be obtained through interaction of string overflow. The subsequent attack scenario is similar to L-1 and L-2.

.7. Details of Attacks from Darpa Engagement

E-1: Information gather and exfiltration. In this attack, the attacker uses the username and password of a target user to log into the system and collect sensitive information, reading */etc/shadow*, */etc/passwd* files, executing commands such as *ifconfig*, *tcpdump*, *psaux*, *groups*, *dirname*, etc., and returning the results to display on */dev/pts/1*.

E-1: Malicious file download and execute. In this attack, the *scp -t./ccleaner* command is used to download a file named *ccleaner* from an unknown address to connect to the external IP 64.95.25.213, then to modify *SrcSinkObject*. The file *ccleaner* is executed to obtain permissions and collect sensitive information from the system. The process also forks a *dbusdaemon* process, which reads files such as */proc/filesystem*, */proc/mount*, and */etc/passwd*. After that, the *ccleaner* and the files generated by it are moved laterally to another host in the intranet through *scp*.

Next a file named *hc* is copied from the external IP 108.25.125.188 to the host, and then the file is executed. In this execution, the file loads a large number of binary and library files, and reads a large number of */proc/directory* files including */proc/net/directory* files and */etc/passwd* and other sensitive files. After that, the file is connected to the external IP, and a file named */tmp/ext96481* is generated. The file is subsequently executed, the permissions were elevated, and the sensitive file */etc/passwd* is read.

E-2: In-memory attack with firefox. In this attack, there is a writable and executable memory space in the *firefox* process. The process also modifies the malicious device file */dev/glx_alsa_675* and increases its permissions. Next, when the *firefox* process communicates with the external IP 86.129.31.201, a malicious library file named */tmp/libnet.so* is dropped. Then the *sshd* process loads the malicious file to connect with the IP address. Subsequently, the malicious *sshd* process extracts multiple files from the host and transfers these files to the at-

tacker's host. Finally, the process creates an executable file named */home/admin/files/docs/audiobackup* and executes it.