

Integrating IDS Alert Correlation and OS-Level Dependency Tracking

Yan Zhai, Peng Ning, Jun Xu

North Carolina State University

Abstract. Intrusion alert correlation techniques correlate alerts into meaningful groups or attack scenarios for the ease to understand by human analysts. However, the performance of correlation is undermined by the imperfectness of intrusion detection techniques. Falsely correlated alerts can be misleading to analysis. This paper presents a practical technique to improve alert correlation by integrating alert correlation techniques with OS-level object dependency tracking. With the support of more detailed and precise information from OS-level event logs, higher accuracy in alert correlation can be achieved. The paper also discusses the application of such integration in improving the accuracy of hypotheses about possibly missed attacks while reducing the complexity of the hypothesizing process. A series of experiments are performed to evaluate the effectiveness of the methods, and the results demonstrate significant improvements on correlation results with the proposed techniques.

1 Introduction

Intrusion detection has received a lot of attention in the past two decades. However, current intrusion detection systems (IDSs) often generate huge numbers of alerts as well as numerous false positives and false negatives. These problems make the reports from IDSs very hard to understand and manage. Many researchers and vendors have proposed various alert correlation techniques (e.g., [6, 7, 15]) to make large numbers of IDS alerts more understandable and at the same time reduce the impact of false positives and false negatives.

Recent alert correlation techniques can be classified into three categories: similarity-based correlation (e.g., [3, 5, 15]), correlation by matching with pre-defined attack scenarios (e.g., [6, 7]), and correlation based on the prerequisite (preconditions) and consequence (postconditions) of individual attacks (e.g., [4, 11]). Each technique has its advantages and disadvantages. However, the correctness of correlation results is strongly affected by the false positives and false negatives among IDS alerts.

Several researchers recently investigated integrating additional information sources into alert correlation to improve its quality. In [10], a formal model named M2D2 was proposed to represent data relevant to alert correlation. The technique in [13] reasons about the relevancy of alerts by fusing alerts with the targets' topology and vulnerabilities, and ranks alerts based on their relationships with critical resources and users' interests. In [16] a statistical reasoning framework is proposed to combine IDS alert correlation with local system state information from tools such as system scanners and system monitors. These approaches can improve the performance of correlation by integrating different sources of security-related information. However, the correlation results are still not yet satisfactory. For example, neither of the approaches can deal with

false negatives quite well. Thus, it is desirable to find additional ways to integrate other information sources to further improve alert correlation.

In this paper, we propose to harness OS-level event logging and dependency tracking to improve the accuracy of alert correlation. OS-level dependency tracking is a recently developed technique to analyze the system operation history toward a given object. It tracks dependency-causing events such as process forking and file operations in the system event log, and spans up a tree of system objects connected by these events from the target object. Though a very useful tool for forensics applications, backtracker has two limitations. Firstly, because it is system call oriented, the complexity of tracking and tracking results can be very high. For example, during a normal system run, the resulting dependency graph of such a tracking (using *Backtracker* [8]) can contain up to tens of thousands of objects and hundreds of thousands of edges. Such complexity with tracking is obviously time and resource consuming, while the tracking results are also hard to understand. Secondly, the tracking is highly dependent on the availability of so-called “detection points”, which are significant evidence of system being attacked. However, such “detection points” are not usually available, making it inconvenient to use in security administration. Integrating event logging and dependency tracking tools with alert correlation can potentially address the above limitations, and at the same time improve the performance of alert correlation.

Our integration method is based on the following observations: Firstly, most attacks have corresponding operations on specific OS-level objects. Secondly, other than a few exceptions, if one attack prepares for another, the later attack’s corresponding operations would be dependent on the earlier one’s. Because logging these system calls is more straightforward than detecting attacks using rules and signatures, such information is considerably more accurate and trustworthy than the IDS alerts. Utilizing such information will improve the performance of alert correlation.

Given the alert correlation results and an OS-level dependency tracking tool, the integration is done in two phases. The first phase is to identify the system objects corresponding to the IDS alerts based on their semantics. The second phase is to verify the relationships demonstrated in the correlation result or discover the missed relationships among the alerts by tracking the dependencies among their corresponding system objects using a dependency tracking tool such as Backtracker.

The contribution of this paper is the development of a framework to integrate the information from OS-level event logging and dependency tracking into IDS alert correlation. With the support of OS-level event logs, we can achieve better accuracy in the final result than the original alert correlation method. We also discuss how such integration can facilitate the hypotheses about possibly missed attacks. Finally, we evaluate the effectiveness of this scheme by performing a series of experiments. Our experiment results show that the integration can greatly improve the correctness of correlation and help making hypotheses about possibly missed attacks. For example, in our experiment, our approach can totally remove the false correlations in all three attack scenarios.

The remainder of this paper is structured as below. Section 2 briefly introduces the background of alert correlation and OS-level dependency tracking. Section 3 discusses the details on how to integrate OS-level object dependency tracking into alert correlation. Section 4 gives experimental results used to evaluate our approach. Section 5 discusses related work. Section 6 concludes and points out some future directions.

2 Background

As discussed earlier, our goal is to improve intrusion alert analysis by integrating OS-level event logging and object dependency tracking into IDS alert correlation. In this section, we give a brief introduction to the alert correlation and OS-level dependency tracking techniques to be used in our method. For alert correlation, we use the method based on attack’s prerequisite and consequence [11], due to the ease to make connections between alert correlation and OS-level objects in this method. The dependency tracking technique used in this paper is *Backtracker* [8], which to our best knowledge is the only such tool available.

2.1 Alert Correlation Based on Prerequisites and Consequences of Attacks

Here we give a brief overview of the alert correlation method in [11]. This method correlates intrusion alerts using the prerequisites and consequences of attacks. Intuitively, the prerequisite of an attack is the necessary condition for the attack to be successful. For example, the existence of a vulnerable service is the prerequisite of a remote buffer overflow attack against the service. The consequence of an attack is the possible outcome of the attack. For example, gaining certain privilege on a remote machine. In a series of attacks, there are usually connections between the consequences of the earlier attacks and the prerequisites of the later ones. Accordingly, we identify the prerequisites and the consequences of attacks, and correlate the detected attacks (i.e., alerts) by matching the consequences of previous alerts to the prerequisites of later ones.

The correlation method uses logical formulas, which are logical combinations of predicates, to represent the prerequisites and consequences of attacks. The correlation model represents the attributes, prerequisites, and consequences of known attacks as alert types. The correlation process is to identify the *prepare-for* relations between alerts, which is done with the help of prerequisite sets and expanded consequence sets of alerts. Given an alert, its prerequisite set is the set of all predicates in the its prerequisite, and its expanded consequence set is the set of all predicates in or implied by its consequence. An earlier alert t_1 *prepares for* a later alert t_2 if the expanded consequence set of t_1 and the prerequisite set of t_2 share some common predicates. An alert correlation graph is used to represent a set of correlated alerts. An alert correlation graph $CG = (N, E)$ is a connected directed acyclic graph, where N is a set of alerts, and for each pair $n_1, n_2 \in N$, there is a directed edge from n_1 to n_2 in E if and only if n_1 prepares for n_2 .

The advantage of this method is that the correlation result is easy to understand and directly reflects the possible attack scenarios. However, as the correlation is solely based on IDS alerts, the result highly depends on the quality of the IDS alerts. For example, the result may contain false correlations when there are false alerts.

2.2 Backtracker

Backtracker is an OS-level dependency tracking tool [8]. Backtracker monitors specific types of OS-level objects, i.e., processes and files. The objects are kept in a log with their properties such as the *uid* of the objects. It also monitors specific dependency-causing system calls like process forking, file reading, and memory sharing, which together are called “high-control events” in [8]. Given the information of a specific object such as the *pid* of a process or the *inode* number of a file, Backtracker identifies the previous objects and system calls that could have potentially affected a target object, and displays

chains of events in a dependency graph. In a Backtracker dependency graph, each node A represents an OS-level object, and each edge $A \rightarrow B$ represents that object B is dependent on object A . Moreover, an edge $A \leftrightarrow B$ is used to represent that objects A and B are potentially dependent on each other. As mentioned earlier, the major limitations of Backtracker are the complexity of its results and the inconvenience to use because of its dependency on the availability of “detection points”. In its later version [9], the tool can also track dependencies between remote hosts by tracking the logged socket ids.

3 Integrating Alert Correlation and OS-Level Dependency Tracking

Our integration is to identify the relevancy between the relationships among IDS alerts and the dependencies among OS-level objects, and then use the OS-level dependencies to verify or discover the relationships among IDS alerts. To identify such relationships, we first look into attacks’ OS-level behaviors.

From the operating system’s point of view, an attack is a set of OS-level events that access or modify a set of system objects. The OS-level objects and operations corresponding to an attack can be derived from the semantics of the attack. In our model, such semantics consist of two parts: one is the prerequisites and consequences of attacks, and the other is the correspondence between the predicates in attacks prerequisites or consequences and the OS-level objects on the host. With such information, we can identify the OS-level objects corresponding to the attacks on the host. For example, given an attack that exploits a vulnerable service as its prerequisite and yields a shell as its consequence, we can identify the corresponding service process and shell process.

Accordingly, the OS-level objects corresponding to an attack can be divided into two sets: the *prerequisite object set*, which are the objects derived from the attack’s prerequisite, and the *consequence object set*, which are the objects derived from the attack’s consequence. These two sets may overlap, because some attacks’ consequences may affect their prerequisite objects. By backtracking among the OS-level objects, we can also find dependencies among those objects at the OS level. Though different from the *prepare-for* relation used in alert correlation, such OS-level dependencies can be utilized to verify or discover the *prepare-for* relations among the alerts.

In our framework, we first extract necessary information from the alerts to identify the corresponding OS-level objects. We then verify the dependencies among alerts by using the OS-level dependencies among their corresponding objects, and thus improve the alert correlation based on the causal relationship. Moreover, by identifying the OS-level objects corresponding to the evidence of possibly missed attacks, and tracking back from those objects, we can improve the performance of existing methods [12, 16] for hypothesizing about possibly missed attacks.

An attack has to have impacts on the local system in order to be observable in the OS-level log. Thus, our method only guarantees improvement of alert correlation for the alerts of successful attacks, though it may provide positive results for some failed attack attempts.

3.1 Identifying OS-Level Objects Corresponding to Intrusion Alerts

Now we discuss how to find the OS-level objects accessed by the attacks which trigger IDS alerts. We call this process the mapping of IDS alerts to OS-level objects. We summarize the semantics carried by an alert that can be used to identify the corresponding

OS-level objects. Firstly, an IDS alert comes with a timestamp, which indicates when the attack happens. Secondly, given an alert, we have the knowledge about how the attack works and how the system should behave in response to it. For example, given a Snort alert “FTP EXPLOIT wu-ftpd 2.6.0”, we know that the corresponding attack exploits a vulnerable wu-ftpd server and forks a root shell. Finally, given local system’s configuration, we can identify which OS-level objects correspond to each predicate in attacks’ prerequisites and consequences. For example, a predicate “Samba server” may correspond to “/usr/sbin/smbd” process on a given computer. Below we discuss how each type of knowledge is used to map the alerts.

Though the number of logged events and objects is large in system logs, the timestamp of each alert can be used to easily narrow down the potentially relevant system objects. In Backtracker’s log, each OS-level object or event is associated with a time period, which is the lifetime of the object or event. (The original Backtracker toolkit does not provide exact timestamp information. We slightly modified Backtracker’s source code to add this functionality.) Given a fixed time period $T = [t_1, t_2]$, an object o can be accessed during T if o ’s lifetime $[t_s, t_e]$ overlaps with T . Given the timestamp of an alert, we can estimate an approximate time window during which all the relevant OS-level activities occur, and then narrow down the scope of OS-level objects that need to be examined. Such a time window has to be relaxed to accommodate delays in OS-level operations and the clock discrepancy between the IDS sensor and the OS.

Given the name of an alert, we have the corresponding attack’s prerequisite and consequence from experts’ knowledge. According to [11], the prerequisite of an attack is a logical combination of predicates, and the consequence of an attack is a set of predicates. Each of those predicates is associated with some OS-level objects such as services, processes, and files. Thus, for each predicate in attacks’ prerequisites and consequences, we can identify the corresponding file or process on the host computer, and represent them as $(predicate, OS-level\ object)$ pairs in the knowledge base. For example, given a pair $(Samba_service(host_IP), "/usr/sbin/smbd")$ in the knowledge base, whenever there is predicate of $Samba_service(host_IP)$, we can locate its corresponding process of “/usr/sbin/smbd”. Thus, after identifying the predicates in an attack’s prerequisite and consequence, we can identify the OS-level objects corresponding to those predicates on the host computer. Based on whether the corresponding predicate belongs to attack’s prerequisite or consequence, those objects can be divided into prerequisite objects and consequence objects.

There are additional constraints for the mapped objects. Firstly, some predicate implies constraints on the properties of its corresponding OS-level objects. For example, the predicate $Root_shell(host_IP)$ implies the privilege of its corresponding OS-level object is root, which is represented by $uid = 0$ for the corresponding object in Backtracker’s log. We represent such a constraint along with the object mapping information in the knowledge base in the form of $(predicate, OS-level\ object, constraint)$. In the above example, we may have a triple $(Root_shell(host_IP), /usr/bin/sh, uid = 0)$ to indicate that the predicate $Root_shell$ is mapped to a process $/usr/bin/sh$, and the process’s uid should be 0.

Secondly, if we focus on successful attacks, we expect to see the prerequisite and consequence of a successful attack at the OS level. In other words, for each predicate

p in an attack’s prerequisite and consequence, there must be at least one object o in the mapped object set corresponding to p , unless the logging tool does not monitor objects corresponding to such predicates. For example, assume the prerequisite and the consequence of a Samba buffer overflow attack are $Vulnerable_Samba_service(dstIP)$ and $Root_shell(dstIP)$, respectively. If this attack is successful, there must be OS-level objects corresponding to these predicates in the object sets. Thus, if such OS-level objects do not exist, the alert must represent a false alert or a failed attack.

Finally, the system activities corresponding to an attack should all be related in Backtracker dependency graph. Moreover, the consequence of an attack should be dependent on the prerequisite of the attack at the OS-level. Thus, for each consequence object corresponding to an attack, it should be a prerequisite object, or dependent on some prerequisite objects of the attack. Thus, we have the dependency constraint: Given alert A and its mapped prerequisite object set P_A and consequence object set C_A , for each consequence object o in C_A , if P_A is not empty, there should exist paths from objects in P_A to o in Backtracker’s dependency graph unless o is also in P_A . If such paths do not exist, the corresponding consequence object o should not be associated with the given alert. For example, given an alert with prerequisite object set $\{service_process\}$ and consequence object set $\{shell_process, file_foo\}$, if $service_process$ is not connected to $shell_process$ in the Backtracker dependency graph, the consequence object $shell_process$ should not be included in the consequence object set.

After mapping IDS alerts to OS-level objects and with OS-level dependency tracking tools, the number of false correlations can be potentially reduced by verifying the dependencies between the corresponding objects. Below we discuss in detail how to identify the OS-level dependencies among alerts with OS-level dependency tracking, and how to use such dependencies to achieve better accuracy in alert correlation.

3.2 OS-Level Dependencies among IDS Alerts

After IDS alerts are mapped to groups of objects within particular time periods in the Backtracker dependency graph, some groups are connected with each other through objects and events in the dependency graph while others are not. It is not difficult to see that an later object is dependent on an earlier object if there exists a path in the dependency graph from the earlier object to the later object. Such paths among these object groups reveal the dependencies between their corresponding alerts. However, such dependencies can be not only malicious attack behaviors but also normal system activities, which makes it different from the prepare-for relations used in alert correlation. To find out the security-relevant dependencies interested by alert correlation, below we discuss in further detail about these dependencies.

As we have mentioned, the OS-level objects corresponding to an alert can be divided into two subsets: the prerequisite object set O_P , which are derived from the attack’s prerequisite, and the consequence object set O_C , which are derived from the attack’s consequence. As discussed earlier, the consequence objects in O_C should be dependent on the prerequisite objects in O_P . Moreover, two alerts being correlated with each other means the earlier attack’s consequence “contributes” to the later attack’s prerequisite. Thus, at the OS level, such a prepare-for relation should be reflected by the paths from the earlier attack’s consequence object set to the later attack’s prerequisite object set (i.e., the later attack’s prerequisite objects are dependent on the earlier attack’s consequence objects). To distinguish such dependencies from other dependencies, we say

alert A is *strongly connected* with alert B if, in the Backtracker dependency graph, there exists a path from one of A 's consequence objects to one of B 's prerequisite objects. Thus, if alert A prepares for alert B , A should be strongly connected with B .

3.3 Verifying the Dependencies among Correlated IDS Alerts

Except for a few types of attacks (e.g., attacks utilizing guest kernels or hidden channels), which are relatively rare as discussed in [8], backtracker is capable of tracking OS-level dependencies among most attacks. In other words, in normal cases, if two alerts are not found strongly connected with each other, there should not be prepare-for relations between them.

Given an alert correlation graph, we can map the alerts to OS-level objects and check whether the correlated alerts are strongly connected. If two correlated alerts are not found strongly connected, the correlation between them is considered false. For example, assume $A \rightarrow B$ are a pair of correlated alerts. To verify this correlation, we first map them into OS-level object sets. If the mappings are successful, there will be corresponding prerequisite and consequence object sets: P_A and C_A of alert A , as well as P_B and C_B of alert B . By tracking back from the objects in P_B with Backtracker, we can verify whether the two alerts are strongly connected. If there does not exist a path between objects in C_A and objects in P_B , we consider the correlation $A \rightarrow B$ false.

Note that two alerts being strongly connected in the dependency graph does not guarantee that the earlier one prepares for the other. This is because OS-level dependencies can be operations of benign programs. That being said, being strongly connected on OS-level dependency graph indicates that the involved attacks have higher possibility to be causally related. Thus, we can use this information to discover attacks missed by IDSs, which lead to missing correlations among alerts.

3.4 Facilitating Hypotheses of Missed Attacks

Integrating IDS alert correlation and OS-level event logging can also help in making hypotheses about possibly missed attacks. Several approaches [12, 16] have been proposed in making hypotheses about possibly missed attacks. Given evidence of attacks being missed, these methods search among known attack types of attacks to fill in the gaps between their correlation graphs and the evidence based on attacks' prerequisites and consequences. However, such search processes can be computationally expensive considering the size of the attack type knowledge base and the number of steps that could have been missed.

Integrating IDS alert correlation with OS-level dependency tracking can facilitate hypothesizing of missed attacks. Signs (evidence) of missing attacks, such as IDS alerts and observed system modifications, can usually be mapped to groups of system objects. Assume a piece of evidence E is mapped to a set of OS-level objects. By tracking backward from those objects, a forest of system objects connected via various events can be generated. For any missed attack, unless it can evade the OS-level dependency tracking, part of its mapped objects must be in this spanned forest. Using the information of the correspondence between predicates and OS-level objects, such a forest of objects can be converted to predicates. Thus, the searching space for possibly missed attacks can be greatly reduced to only the attacks related to the predicates.

For example, an attacker takes the following attack steps: (1) Launch a buffer overflow attack toward a vulnerable Samba server, which yields a root shell; (2) Delete the

web page files via the shell. Assume all those activities are missed by the IDS, while the file deletion is detected by some file system integrity monitoring tool, which is taken as evidence indicating previous attacks being missed. By tracking back from the deleted file, the file is found dependent on the following objects in the OS-level event log: a smbd process and a shell process forked by the smbd process. Thus, when searching for possibly missed attacks, we can limit the search within attacks related to Samba and shell. Since only Samba is an initial system service, we hypothesize there is a Samba-related attack missed. By trying to map each candidate attack to OS-level objects, we can eliminate the majority of invalid hypotheses. The uncertainty within the remaining results is affected by the knowledge we have about the local system (e.g., the version of Samba) and the attacks (e.g., the number of attack types in the knowledge base).

4 Experimental Results

We have performed a series of experiments to validate the effectiveness of our method. Since our method requires actual system responses to real attacks, we developed three attack scenarios in our lab, in which an attacker launches a sequence of attacks against a computer monitored by both the Backtracker and the Snort IDS.

The Backtracker was configured on the target server and slightly modified to log system calls with their timestamps. Two vulnerable services were running on the server: Samba 2.2.8 and icecast 1.3.11. Snort 2.40 was installed on the server to monitor the network traffic as an IDS sensor. To detect more attacks, we used the “Bleeding Snort Rulesets” [1] with Snort. Clock drifting is not considered in our experiments because both programs were running on the same computer. We also injected background traffic during the experiments to mimic an operational network. The background traffic was collected on the target machine when it was connected to our campus network, and was manually verified to contain no attacks toward the target machine. We also injected some failed attempts of wu-ftp buffer overflow attacks into the background traffic.

Due to space reasons, we only discuss in detail the experimental results on Scenario 1 here in this paper. Description of the other two scenarios and additional information about these experiments can be found in the full version of this paper [17].

4.1 Details of Scenario 1

Our first attack scenario exploits a vulnerable Samba server. It includes the following attack steps:

1. Two remote buffer overflow attack attempts exploiting the vulnerable Samba server (resulting 2 remote root shells).
2. Uploading and starting a server daemon of a DDoS tool called TFN (Tribe Flood Network) on the target host.
3. Using TFN’s client program to direct the TFN server daemon on the victim server to start SYN flood and UDP flood attacks against another computer.

The above attacks took about 5 minutes. During the period, Backtracker logged 81,613 events. Moreover, the Snort sensor raised 9 “NETBIOS SMB trans2open buffer overflow attempt” (“smb-bof”) alerts (No. 1–8 and No.10), 15 “DDOS tfn2k icmp possible communication” (“tfn2k-icmp”) alerts (No. 12–26), and 2 “ATTACK-RESPONSES

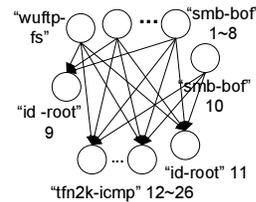


Fig. 1. Original Correlation Graph

id check returned root” (“id-root”) alerts (No.9 and No.11). The background traffic triggered 32 alerts related to the target server, which include 8 “SCAN nmap TCP” (“nmap-tcp”) alerts, 23 “SNMP public access udp” (“snmp-udp”) alerts, and 1 “FTP EXPLOIT wu-ftpd 2.6.0 site exec format string overflow Linux” (“wuftp-fs”) alert. Among the 3 types of alerts, the third one is triggered by the failed attempt of wu-ftpd buffer overflow attack injected into the background traffic.

Table 1. OS-level Objects Corresponding to the Alerts in Scenario 1

Alert	Prerequisite Objects	Consequence Objects
“smb-bof” 8	{smbd_2717}	{sh_2720}
“smb-bof” 10	{smbd_2717}	{sh_2725}
“id-root” 9	{sh_2722, /usr/bin/id_324551}	Null
“id-root” 11	{sh_2727, /usr/bin/id_324551}	Null
“tfn2k-icmp” 12 26	{td_2737}	Null

Using the alert correlation method proposed in [11], we generated the correlation graph shown in Fig. 1. Obviously, it contains many false correlations due to the false positives within the alerts. Using the Backtracker’s log and the semantics of these alerts, we mapped these alerts to a number of OS-level objects, as listed in Table 1.

For each alert prepared by other alerts in Fig. 1, we generated Backtracker dependency graphs by tracking back from their prerequisite objects. In these graphs, we look for paths from an earlier alerts’ consequence objects to a later alert’s prerequisite objects if the former prepares for the later. Examples of such paths found in our experiments are shown in Fig.2(a). According to our previous discussion, when such a path exists, the two alerts are strongly connected and thus the correlations between them are verified at the OS level. Otherwise, the correlation would be considered false. In this way, we can verify each of the correlations in the original correlation graph, remove those that are verified to be false, and finally come up with a new correlation graph. The new correlation graph for scenario 1 is shown in Fig. 3(a). We can see it is the correct correlation graph of the reported Snort alerts based on the actual attack scenario.

In the correlation graph, the “smb-bof” alerts prepare for the “tfn2k-icmp” alerts because the consequence of the former (i.e., root_shell (dstIP)) implies the prerequisite of the later (i.e., tfn2k_server_daemon ()). This implication indicates that some activities are missed between the two alerts.

By backtracking from the prerequisite object set {td_2737} of the tfn2k alert, a tree of OS-level objects are spanned. Because the consequence object set {sh_2720} of the “smb-bof” alert is among them, we focus on the paths linking the two object sets. Along the path shown in Fig. 2(c), we find the following OS-level objects being involved: process object “ftp_2732”, file ob-

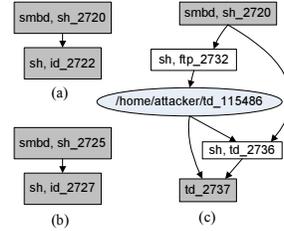


Fig. 2. Paths Between Alerts’ Mapped Objects

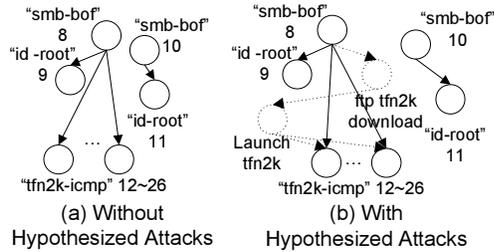


Fig. 3. New Correlation Graphs

ject “/home/attacker/td_115486”, and process object “td_2736”. Thus, we can limit the searching within the activities related to the ftp and tfn2k server program. It is easy to see that the attacker downloaded the tfn2k server program via ftp and launched it via the shell. These hypotheses are shown in dotted circles and lines in Fig. 3(b).

4.2 Overall Evaluation

Assume an alert correlation method outputs that an alert A prepares for another alert B . If both alerts are detections of actual attacks, and the attack corresponding to alert A is indeed used to prepare for the later attack corresponding to alert B , we consider this correlation as a *true correlation*. Otherwise, it is considered a *false correlation*. Moreover, if one attack is used to prepare for another attack, but there is no correlation corresponding to these attacks (due to missing detection or incorrect correlation), we say there is a *missing correlation*. In our experiments, since we know the details of the attack scenarios, we can easily identify true, false, and missing correlations.

We use two metrics, false correlation rate and missing correlation rate, to evaluate the overall performance of alert correlation before and after integrating Backtracker’s results. Given a set of correlated alerts, the *false correlation rate* is the ratio between the number of false correlations over the total number of correlations generated by alert correlation. The *missing correlation rate* is the ratio between the number of missing correlations over the total number of correlations between *attacks*. Intuitively, these two metrics show the degrees of correctness and completeness of the correlation. Obviously, the smaller these two metrics are, the better performance alert correlation has.

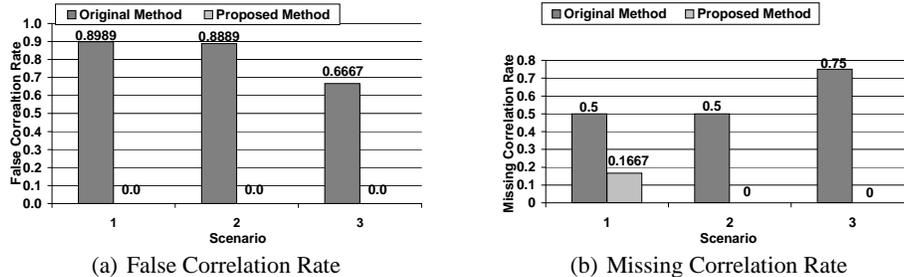


Fig. 4. Compare between the Original Correlation Method [11] and the Proposed Method

Fig. 4(a) shows the false correlation rates in all three attack scenarios. As we can see, our proposed method reduces the false correlation rate significantly in all three scenarios. Indeed, false correlations are completely removed in all scenarios. This is not surprising, because OS-level dependency provides another way to properly verify the correlation between alerts through trustworthy information kept in OS-level logs.

Fig. 4(b) shows the missing correlation rates in the three scenarios. We can see significant reduction in missing correlation rate in all three scenarios. While the missing correlation rate is reduced to 0 in scenarios 2 and 3, the missing correlation rate of the first scenario is still non-zero after making hypotheses. This is because the DDoS attack (via the tfn server) is neither detected by Snort, nor hypothesized by our approach.

Our experiments confirmed that OS-level object dependency graphs can often be too complicated to understand in reality. For example, during the 30 minutes’ period of the third scenario, Backtracker logged more than 410,000 events, resulting a dependency graph of more than 4,000 nodes. Manually analyzing such a complicated graph

requires lots of time and very detailed knowledge about attacks' OS-level behaviors. Our proposed method solves this problem because: 1. our method tries to *verify* alert correlations instead of to *detect* attacks, only moderate information about attacks' OS-level behaviors is required, and 2. the verification processes can be done automatically by computer programs instead of human experts.

The experiment results also showed how the proposed method can help make hypotheses about possibly missed attacks. In the attack scenarios, one type of attacks are hard to be detected by any IDS, which are the attackers' "legitimate" activities after the break-in. It is very difficult to guess about such activities without the OS-level dependency information.

5 Related Work

Our technique is closely related to the alert correlation techniques based on prerequisites (pre-conditions) and consequences (post-conditions) of attacks proposed in [4, 11, 14] as well as OS-level dependency tracking techniques proposed in [8, 9], which have been discussed in the Introduction and Section 2. The work closest to ours is [9]. In [9], King et al. proposed to use the dependency information from Backtracker to study attacks within and across the hosts, as well as prospected the application of Backtracker in alert correlation. Although they demonstrated the potential benefit of combining backtracker analysis and IDS alert analysis with some heuristic case studies, the main focus of [9] is the backtracking among remote hosts. It does not give any specific method for combining the two different analyses. All the analyses in *EnrichAlerts* are done manually, which is computationally expensive and difficult in practice. In our approach, after discussing the difference and resemblance between the OS-level dependencies and the dependencies among attacks, we proposed a specific method to automatically combine the analyses.

Several techniques also use local system information to reason about the causal relationships between IDS alerts. In [16], local system state information is used to reason about the correlation of alerts. In [2], such information is used to analyze the vulnerability of the system and study the potential attacks that could compromise the system. These techniques are complementary to the approach proposed in this paper.

Some approaches have been proposed in [12, 16] to make hypotheses about attacks possibly missed by IDSs. [12] makes hypotheses based on the prerequisite and consequences of attack types when similarities are found between alerts in separate correlation graphs. [16] uses similar attack type knowledge to make hypotheses upon inconsistencies between observed facts and alert correlation graph. Our techniques can facilitate the hypothesis process by bringing additional information from OS-level dependency tracking, thus to reduce the search space for the possibly missed attacks.

6 Conclusion and Future work

In this paper, we developed a series of techniques to integrate the alert correlation method (based on prerequisites and consequences of attacks) and OS-level object dependency tracking. A critical step in this integration is to map IDS alerts to OS-level objects, so that connections between alert correlation and OS-level objects can be established. We also identified a number of constraints that the OS-level objects should satisfy if they are relevant to the IDS alerts (or attacks) that are correlated. By using these

constraints, we can verify the IDS alerts as well as the correlation between IDS alerts, and reduce false correlations. Moreover, the dependency between OS-level objects can also facilitate the hypotheses of attacks possibly missed by the IDSs. Our experimental evaluation gave favorable results, showing that OS-level dependency tracking can significantly reduce false correlations when integrated with the alert correlation method.

Several issues are worth doing in our future research. First, we will investigate more techniques to integrate OS-level dependency tracking with techniques to hypothesize about possibly missed attacks. Second, we want to investigate the possibility of learning “normal” patterns of dependencies among OS-level objects from training period to improve the verification of correlations and making hypotheses.

References

1. Bleedingsnort. www.bleedingsnort.com.
2. P. Ammann, D. Wijesekera, and S. Kaushik. Scalable, graph-based network vulnerability analysis. In *Proceedings of the 9th ACM CCS*, 2002.
3. F. Cuppens. Managing alerts in a multi-intrusion detection environment. In *Proceedings of the 17th Annual Computer Security Applications Conference*, December 2001.
4. F. Cuppens and A. Mieke. Alert correlation in a cooperative intrusion detection framework. In *Proceedings of the 2002 IEEE Symposium on Security and Privacy*, May 2002.
5. O. Dain and R. Cunningham. Building scenarios from a heterogeneous alert stream. In *Proceedings of the 2001 IEEE Workshop on Information Assurance and Security*, June 2001.
6. O. Dain and R. Cunningham. Fusing a heterogeneous alert stream into scenarios. In *Proceedings of the 2001 ACM Workshop on Data Mining for Security Applications*, Nov. 2001.
7. H. Debar and A. Wespi. Aggregation and correlation of intrusion-detection alerts. In *Recent Advances in Intrusion Detection*, LNCS 2212, pages 85 – 103, 2001.
8. S. King and P. Chen. Backtracking intrusions. In *Proceedings of the 2003 Symposium on Operating Systems Principles (SOSP)*, October 2003.
9. S. King, Z. Mao, D. Lucchetti, and P. Chen. Enriching intrusion alerts through multi-host causality. In *Proceedings of the 12th NDSS*, 2005.
10. B. Morin, L. Mé, H. Debar, and M. Ducassé. M2D2: A formal data model for IDS alert correlation. In *Proceedings of RAID 2002*.
11. P. Ning, Y. Cui, and D. S. Reeves. Constructing attack scenarios through correlation of intrusion alerts. In *Proceedings of the 9th ACM CCS*, 2002.
12. P. Ning, D. Xu, C. Healey, and R. St. Amant. Building attack scenarios through integration of complementary alert correlation methods. In *Proceedings of the 11th Annual Network and Distributed System Security Symposium (NDSS '04)*, pages 97–111, February 2004.
13. P. Porras, M. Fong, and A. Valdes. A mission-impact-based approach to INFOSEC alarm correlation. In *Proceedings of RAID 2002*.
14. S. Templeton and K. Levitt. A requires/provides model for computer attacks. In *Proceedings of New Security Paradigms Workshop*, pages 31 – 38. ACM Press, September 2000.
15. A. Valdes and K. Skinner. Probabilistic alert correlation. In *Proceedings of the 4th International Symposium on Recent Advances in Intrusion Detection*, pages 54–68, 2001.
16. Y. Zhai, P. Ning, P. Iyer, and D. Reeves. Reasoning about complementary intrusion evidence. In *Proceedings of the 20th Annual Computer Security Applications Conference*, Dec. 2004.
17. Y. Zhai, P. Ning, and J. Xu. Integrating IDS alert correlation and OS-level dependency tracking. Technical Report TR-2005-27, Department of Computer Science, North Carolina State University, 2005.