Insider Threat Control: Using a SIEM signature to detect potential precursors to IT Sabotage

CERT Insider Threat Center

April 2011

Software Engineering Institute | Carnegie Mellon

Table of Contents

Introduction

This paper describes the development and proposed application of a Security Information and Event Management (SIEM) signature to detect possible malicious insider activity leading to IT sabotage. In the absence of a uniform, standardized event logging format, this paper presents the signature in two of the most visible public formats, Common Event Framework (CEF) and Common Event Expression (CEE). Because of the limitations of these formats, the SIEM described in this paper employs an operational version of the proposed signature in an ArcSight environment.[1]

Insider Threat Database

The CERT Insider Threat Center database currently contains over 550 cases of actual malicious insider crimes. Our research has revealed that most crimes fit one of four categories:

- IT Sabotage
- Theft of Intellectual Property
- Fraud
- Espionage

We define a malicious insider as a current or former employee, contractor, or business partner who

- has or had authorized access to an organization's network, system or data
- intentionally exceeded or misused that access
- negatively affected the confidentiality, integrity, or availability of the organization's information or information systems

This paper focuses on the 123 cases categorized as IT sabotage. Insider IT sabotage is defined as an insider's use of information technology to direct specific harm at an organization or an individual. Each entry contains general details about the case, including a timeline of the incident. The specific codebook items we focus on, which were derived from our analysis of the database, are the attack location (i.e., on site vs. remote access) and attack time (i.e., during work hours vs. outside working hours), as well as what type of protocol the malicious insider used for remote access (i.e., SSH vs. Telnet vs. RDP).

The purpose of this analysis is to develop an SIEM signature to detect the presence of a malicious insider based on key indicators related to IT sabotage activity.

---

[1] Note: the use of ArcSight by the authors of this paper is not intended as an endorsement of the ArcSight platform, but rather reflects the availability of this platform to the authors.

A Note on Signature Application

Technical signatures developed by the CERT® Insider Threat Center are generally designed to be applied towards a particular user or group of users. These signatures are not intended to be applied to all users across the enterprise, as doing so will generate a large number of false positives. The cases in our database reveal that almost all insiders involved in acts of IT Sabotage displayed behavioral indicators prior to committing their crimes. The respective organizations could have, and should have, acted upon these behavioral precursors to prevent the crimes from taking place. Examples of such behavioral indicators include but are not limited to

- Conflicts with co-workers or supervisor
- improper use of organization information assets
- Sanctions
- Rule violations and/or security violations

Prior to applying this signature, the organization should facilitate proper communication and coordination between relevant departments across the enterprise, especially information technology, information security, human resources, physical security, and legal. This cooperation is necessary to ensure that any measures taken by the organization to combat insider threat comply with all organizational, local, and national laws and regulations. Once users are identified who warrant targeted monitoring via this signature, the organization will then be able to determine the appropriate user names, account names, host names, and/or host addresses to enter into the signature to make the alert volume more meaningful and manageable. Finally, it is critical that the organization create a clearly defined policy for targeted employee monitoring, which is consistently enforced.

Database Analysis

We conducted a brief analysis on the IT Sabotage cases in the database based on the following questions to find what information can be used to develop a SIEM signature:

1. What time did they attack? After hours vs. business hours?
2. How many insiders attacked using virtual private network (VPN) vs. in the office?
3. What protocols do insiders use for remote connection? Secure Shell Host (SSH), Telnet, Remote Desktop Protocol (RDP)?

Regarding the time of attack, based on the 123 cases of IT sabotage, we found that 26 percent of the attacks occurred during work hours, 35 percent outside of working hours, and in 39 percent of the cases the time of attack was unknown. Figure 1 graphically represents these findings. Breaking this down further, out of the cases for which the time of attack is specified, 58 percent of the attacks occurred outside normal working hours and 42 percent of the attacks occurred during work hours.

Figure 1: Time of Attacks for IT Sabotage Cases

Another significant finding concerned the number of insiders who attacked using VPN vs. the number of insiders who attacked while in the office. We found that 54 percent of the attacks used remote access and 27 percent of the cases occurred on site. Only 19 percent of the cases did not specify the location of the attack. Therefore, if we again only consider those cases that specify the location of the attack, 66 percent of the attacks occurred using remote access and 34 percent of the attacks occurred on site. Figure 2 presents these findings.


Figure 2: Location of Attacks for IT Sabotage Cases

Note that the VPN connection by itself does not indicate malicious activity. Rather, the insiders most often used a remote connection to the target system after they established a VPN connection

with the victim network. For this reason, we do not include any VPN traffic as a monitored protocol, but rather, we do include the VPN username in cases where that account may differ from the user's regular username. Figure 3 denotes the typical sequence of events associated with a remote attack via VPN.


Figure 3: Typical remote attack activity via VPN

The specific protocols insiders use for remote connections is not currently coded in the database. However, through interviews with some of the actual perpetrators, as well as through a more detailed analysis of these cases, we discovered that the most common known ports used for remote attacks are port 22 (SSH), 23 (Telnet) and 3389 (Terminal Services, or RDP). Since a majority of malicious insiders used remote access for their attacks, we considered instances of connections to these three ports as suspicious in the development of our signature. Each organization will need to account for other protocols used in their own environments to make sure they are monitoring all possible channels of communication.

Based on this analysis of the database, we found that the relevant indicators to be included in this particular control are the location of the attack and the time of the attack. Also, as previously mentioned, since remote access is a common method of attack, it is important to consider the type of protocol the attacker uses (although this information was not specifically coded in the database). This information is the basis for our SIEM signature.

SIEM Signature

Note that this signature is to be applied only to individuals who warrant increased scrutiny. **This signature should not be applied to all privileged users as it will generate inordinate false positives.**

The characteristics of the attacker involve someone accessing the organization's information systems remotely, outside normal working hours. With these characteristics, we developed the following signature:

```
Detect <username> and/or <VPN account name> and/or <hostname> using
<ssh> and/or <telnet> and/or <RDP> from <5:00 PM> to <9:00 AM>
```

The purpose of the signature is to detect the identity of the attacker, what remote connection protocol he or she is using, and whether this activity is occurring outside normal working hours. The identity of the attacker can be retrieved through any or all of the following parameters: username, VPN account name, or hostname. Similarly, the remote connection protocol can be any or all of the following: SSH, Telnet, or RDP. We have based the signature on the following key fields: username, VPN account name (in case this account name is different from the local username), hostname of the attacker, and whether they are using SSH, Telnet, or RDP.

With this basic structure in mind, two standards we used for creating the signature were the Common Event Format, developed by ArcSight, and the Common Event Expression, developed by MITRE. Brief summaries of each standard are provided in the following sections.

Common Event Format

The Common Event Format is an even interoperability standard developed by ArcSight. The purpose of this standard is to improve the interoperability of infrastructure devices by instituting a common log output format for different technology vendors. It assures that an event and its semantics contain all necessary information. CEF is an extensible, text-based format designed to support multiple device types in the easiest way possible. It defines syntax for log records consisting of a standard header and a variable extension that is formatted as key-value pairs. This format contains the most relevant information, which makes it easier for event consumers to parse and use them. The format of CEF is: (**header**/**extension**):

```
CEF:Version|Device Vendor|Device Product|Device
Version|Signature ID|Name|Severity|Extension
```

The *Version* identifies the version of the CEF format. The *Device Vendor*, *Product,* and *Version* uniquely identify the type of sending device. The *Signature ID* identifies the type of event reported. The *Name* represents a human-readable and understandable description of the event, and the *Severity* reflects the importance of the event. The *Extension* part of the message is a placeholder for additional fields, which are part of a predetermined set.

Using this standard and the key indicators we identified during the database analysis, we developed the following two CEF-based, SIEM signatures to identify suspected attackers. The first signature is for Microsoft products. It identifies a suspected attacker by logging his or her username and hostname.

```
CEF:0|microsoft|activedirectory|2011|001|username logged
in|10|suser=<username> src=<10.0.0.1> shost=<hostname>
```

With this information, the second signature is for Snort products. It identifies an attacker who initiates a remote connection to TCP port 22, 23, or 3389. It logs the username/IP address of the suspected attacker gathered from the first signature, the destination address, or the source hostname.

```
CEF:0|sourcefire|snort|2.9|002|remote connection from <suser> or <src>
or <shost> to <dst>|src=<10.0.0.1> or shost=<hostname> prot=TCP
dpt=<22,23,3389> start=<17:00:00> end=<08:00:00>
```

Since a single CEF cannot be used to draw from two separate products, these two signatures are used together to identify a suspected malicious insider.

Common Event Expression

The Common Event Expression (CEE) Architecture defines an open and practical event log standard developed by MITRE. Like CEF, the purpose of CEE is to improve the audit process and users' ability to effectively interpret and analyze event log and audit data. It also enables the creation of useful and efficient log records within applications. It standardizes the event-log relationship by normalizing the way events are recorded, shared and interpreted.

The basic components of CEE are dictionary and event taxonomy, logging recommendations, log syntax, and log transport. Event records are guided by log recommendations (suggested events to log). Log messages are exchanged via a common log transport (standard communications mechanisms, such as XML, SMTP, Syslog, etc.); log messages are received in common log syntax (consistent data elements and format, such as XML); and the dictionary and event taxonomy specifies the event in a common representation (standard field names, terminology, and event types, such as a user login, service restart, or network connection).

Using the CEE format, we developed a signature based on the key indicators of insider IT sabotage. A sample signature using arbitrary data for <logTime>, <user>, <src>, and <shost> is detailed below in XML format.

```
<event name="remote connection by suspected malicious insider">
      <logTime>2011-03-17T12:17:32</logtime>
      <suser>maliciousinsider</suser>
      <src>10.0.0.1</src>
      <shost>insider_system</shost>
      <prot>TCP</prot>
      <dpt>{22,23,3389}</dpt>
      <start>17:00:00</start>
      <end>08:00:00</end>
</event>
```

The signature identifies a suspected attacker who is using a remote connection to log onto the organization's internal system using TCP port 22, 23, or 3389 outside normal working hours. It also logs the time the event was recorded.

Applying the Signature

In order to implement and test the signature in a production environment, we deployed it in an ArcSight platform. The ArcSight signature developed and fully tested using our key indicators is detailed below.

```
 ((Attacker User Name = <username> OR Attacker Host Name = <hostname>)
 AND (Target Port = 3389 OR Target Port = 23 OR Target Port = 22) AND
 Manager Receipt Time Between (17:00:00.000,08:00:00.000) AND Target
 Address = any)
```

This signature generates an alert if, after normal working hours (i.e., between 5:00p.m. – 8:00a.m.), an attacker is connected to any machine via port 3389 (RDP), 23 (Telnet) or 22 (SSH). To identify the attacker, the signature logs the attacker's username or hostname.

Note that the major determinant of utility and success of this signature is proper identification of users to whom this signature will be applied. **This signature should not be applied to a general user population as it will generate a large number of false positives**. Privileged users, such as systems administrators, typically connect remotely to various systems outside office hours in the normal course of their daily activities. To determine which users merit more targeted monitoring through this signature, the organization will have to rely on human resources records to properly identify employees who have exhibited concerning behaviors that warrant closer inspection.

Example Case Study

A case example in which this signature could have been used involved an insider who was formerly employed as a software engineer by the victim organization, a high-technology company. The insider was responsible for managing an automated manufacturing system. During the work week, the insider maintained a constant remote access connection from his home to the organization's network.

The insider, who had previously worked in another department at the organization, was terminated due to poor performance. Prior to informing the insider of his termination, the organization terminated the insider's network access, but failed to check if the insider's remote access connection was active.

The incident occurred the day after the insider's termination, outside of working hours. While under the influence of alcohol, the insider used the open VPN connection to remotely connect to critical systems and shut down the organization's manufacturing system by deleting critical files.

Due to the insider's actions, the organization lost four hours of manufacturing time and had to load backup data to restart the manufacturing process. The incident cost the organization $20,000 to remedy. Connection and activity logs connected the insider to the incident, and the insider was arrested and convicted.

In this case, since the insider remotely accessed the organization's information systems outside of working hours using the insider's own account, the signature we developed would have alerted on this suspicious activity even before the insider sabotaged the data. The signature would have notified system administrators to the insider's initial VPN connection Monday evenings and every day during the week, since the insider left it connected all day, all week. It would have logged from whose account and from where the connection was coming, and could have potentially detected the insider before he deleted the organization's critical information. The signature would also have alerted on the insider's remote connectivity to the critical systems and his deletion of operational files. Without a signature like this, the insider was able to exploit the vulnerability the organization created by failing to disable the insider's connections upon termination.

Conclusion

Organizations should require information security personnel to regularly communicate with different departments across the enterprise, especially with HR and legal, to determine if any employees have exhibited concerning behavior that may warrant targeted monitoring. In CERT's Insider Threat database, the vast majority of insiders committing IT sabotage were guilty of policy violations, and in some cases criminal activity, prior to the execution of their attack. In most cases, insiders carried out their attack via a VPN connection, from which they launched remote connections to their target systems. Organizations should identify suspicious insiders and apply the SIEM signature described in this paper to ensure that their actions are closely monitored.