

ALERT-ID: Analyze Logs of the Network Element in Real Time for Intrusion Detection

Jie Chu¹, Zihui Ge², Richard Huber³, Ping Ji¹,
Jennifer Yates², and Yung-Chao Yu³

¹ Graduate Center CUNY, 365 Fifth Avenue, New York NY 10016

² AT&T Labs - Research, 180 Park Avenue, Florham Park NJ 07932

³ AT&T Labs, 200 Laurel Avenue, Middletown, NJ 07748

Abstract. The security of the networking infrastructure (e.g., routers and switches) in large scale enterprise or Internet service provider (ISP) networks is mainly achieved through mechanisms such as access control lists (ACLs) at the edge of the network and deployment of centralized AAA (authentication, authorization and accounting) systems governing all access to network devices. However, a misconfigured edge router or a compromised user account may put the entire network at risk. In this paper, we propose enhancing existing security measures with an intrusion detection system overseeing all network management activities. We analyze device access logs collected via the AAA system, particularly TACACS+, in a global tier-1 ISP network and extract features that can be used to distinguish normal operational activities from rogue/anomalous ones. Based on our analyses, we develop a real-time intrusion detection system that constructs normal behavior models with respect to device access patterns and the configuration and control activities of individual accounts from their long-term historical logs and alerts in real-time when usage deviates from the models. Our evaluation shows that this system effectively identifies potential intrusions and misuses with an acceptable level of overall alarm rate.

1 Introduction

A fundamental aspect of network security is securing the networking infrastructure itself, which can be particularly challenging in a large scale enterprise or ISP (Internet service provider) network. In such networks, hundreds or thousands of routers and switches are widely dispersed among a geographically diverse set of offices and are typically managed by a large team of network operators. It is imperative that the networking infrastructure and the information contained therein be fully protected against any malicious priers and attackers. For example, information available at networking devices, such as router configuration and traffic statistics, may contain confidential business data of tremendous value to a business competitor. Divulging such information will likely result in a significant disadvantage to the ISP's business. Leakage of some critical security information in the router configuration such as QoS policy or firewall/ACL (Access Control List) settings may subject the network to crafted and targeted attacks such as DDoS (Distributed Denial of Service) attacks. Or in an even more devastating scenario, malicious attackers gaining privileged access to the networking device might

alter the network configuration to create havoc and paralyze the entire network and the services it supports.

Given the risk of severe consequences, large scale networks typically devise and deploy a range of security and protection measures for their networking devices. One common practice is to utilize a combination of *periphery protection* and centralized *authentication and authorization* for communication to networking devices. By restricting premises access, unauthorized persons are blocked from gaining physical access to networking devices. Through careful configuration of ACLs at all network edge routers, unauthorized network traffic is also blocked from reaching network devices. And finally, technologies such as TACACS+ (Terminal Access Controller Access-Control System Plus) [2] and RADIUS (Remote Authentication Dial In User Service) [12], ensure that only authenticated users (i.e. authorized network operators/administrators) have access to routers and switches (either directly or remotely over the network).

The architecture above is very effective against threats from external attackers when working properly. However, there is always the possibility that building security is breached, allowing physical access to router hardware, or that ACLs on an edge router are misconfigured, admitting attacking traffic. Furthermore, with a large team of network operators, compromised users or compromised user accounts can be a critical source of potential security troubles arising inside the network.

In this paper, we propose to add another layer of defense for networking infrastructure by overseeing *all* operations being done in the network, and automatically detecting and raising alarms for “suspicious” activities. We leverage the existing authentication and authorization framework and collect router/switch access logs in real-time. We develop an anomaly detection system that compares on-going router/switch access activities against a set of patterns or profiles constructed from historical data, and once an anomaly is identified, triggers an alarm to network security managers for further investigation of potential intrusions and misuses.

Although the concept of intrusion detection system is well established in computer system security, applying the idea in networking device management remains unexplored, interesting, and challenging. To detect abnormal activities, we must obtain data on routine/normal network management activities in a large scale network, analyze that data, and determine what features best distinguish normal activities from abnormal ones. In our study, we base our analysis on actual network data from one of the largest ISP networks, which comprises tens of thousands of routers distributed worldwide. We conduct an in-depth analysis on a wide range of different characteristics about operators’ access patterns and identify useful features. The effectiveness of an intrusion detection system is known to be limited by noisy baseline behavior and hence high false positives. Thus, when developing the detection methodology and the prototype system for capturing potential intrusions and misuses, we focus on managing false positives to be well within an acceptable range. Any given attack is likely to come from a small number of source subnets or accounts. Thus we aggregate detected “threat scores” by their origin source addresses and login accounts. This allows us to amplify the signal of offense and hence be able to detect offenders while they are still exploring the network before large-scale damage is inflicted.

Our contribution in this paper can be summarized as follows:

- We propose to systematically monitor and analyze the networking device access logs to protect the networking infrastructure. To the best of our knowledge, this is the first study that focuses on monitoring and auditing networking device access and control logs to catch anomalous activities.
- We conduct an in-depth analysis and characterization on the TACACS+ logs collected over more than six months from a tier-1 ISP network. We further identify a set of features that can be utilized to distinguish suspicious activities from normal operations — such as the login ID and origin IP prefix association pattern, the daily number of distinct routers accessed, and the number of hops over which an operator logs on to a router from a different router.
- we develop a tool for the ISP network. Our controlled experiment shows that it successfully identifies injected “malicious” activities – with corresponding threat scores significantly higher than those of day-to-day operational activities. When used in real operation, the system produces no more than a few threat alerts per day – a level at which network security managers are comfortable in conducting further investigation. Many of the threats detected have been found interesting and worth examining.

The rest of the paper is organized as follows. In Section 2, we provide an overview of operational management activities in large scale IP networks and a brief introduction of the authentication, authorization, and accounting system from which we collect logs. Section 3 presents our analysis result on the characteristics of normal operation activities. Section 4 describes the rules and detection system that we build for detecting and alerting on suspicious router accesses and controls. We evaluate our overall system performance in Section 5. Based on our operational experience, we propose a further enhancement to the system and evaluate its effectiveness in Section 6. We discuss related work in Section 7 and finally conclude the study in Section 8.

2 Background

2.1 Managing IP Networks

We first provide an overview of the various types of management activities in large scale IP networks. We describe these in the setting of a global ISP network although many of them are fundamental to large enterprise networks or regional ISP networks as well.

Managing a global ISP network requires a large team of network operators. These operators are typically organized in tiers – lower tier operators respond to routine issues following a set of predefined standard procedures, while more complex matters are escalated to upper tier operators, who have deeper knowledge and understanding of the network. Extremely complicated issues are escalated to a small group of experts, possibly including designers and support teams of vendors of involved devices.

Different tiers of operators have different functional roles. Some may be dedicated to the care of a high profile enterprise customer, in which case they will frequently access provider edge (PE) routers but seldom touch backbone routers. Some operators may

be responsible for servicing the metropolitan area network for a certain region. Others may oversee control plane health (e.g., router CPU utilization) for the entire network. Depending on their role, operators are expected to have distinct patterns of network management activities.

Network operators often exercise control over routers and switches by logging on to the device. Today, nearly all networking devices support console access via direct connection to them and remote access via `ssh/telnet`. Control is exercised by invoking a sequence of commands through the Command Line Interface (CLI) of the device's operating system. For example, on Cisco IOS, typing

```
ping 1.2.3.4
```

triggers a ping test from the router to the IP address. And typing

```
enable
configure terminal
interface Ethernet0
shutdown
exit
```

administratively shuts down the interface *Ethernet0* at the router.

Note that Cisco IOS supports two different access levels – user level and privileged level. The `enable` command in the above example enters the privileged level, in which configuration change (`configure terminal`) is allowed. Such capability is widely supported on other vendor systems such as Juniper JunOS as well. In addition, AAA systems (described below) support finer grained command groups. A user cannot invoke commands outside of his/her predetermined access levels or command groups.

In addition to operators typing commands via the CLI, ISPs rely on a broad range of automated tools for their network management activities. These tools are typically designed to perform a specific set of functions. For example, automation tools/systems that perform configuration auditing periodically sweep through the entire network issuing a `show running-config` command to collect active router configurations. Another tool might collect hardware, traffic, or protocol status and statistics information by logging into the routers of interest and invoking commands such as `show process CPU history`, `show interfaces POS 1/0`, and `show ip bgp summary`. The tools may use designated logins when requesting access to networking devices.

Using the combination of function-level controls via various automated systems and manual command-level controls, operators are able to accomplish a wide range of network management tasks including provisioning and decommissioning customer services, troubleshooting networking and service problems, performing device life cycle management, taking measurements, and monitoring the health of the network and services.

2.2 Authentication, Authorization and Accounting

The networking infrastructure in large scale networks is typically protected by an AAA (Authentication, Authorization, and Accounting) system. There are two mainstream AAA frameworks widely used commercially – TACACS+ (Terminal Access Controller Access-Control System Plus) [2] and RADIUS (Remote Authentication Dial In User Service) [12]. While differing in some specifics, such as whether authentication and authorization are separately maintained in user profiles, both systems use one or more

common servers to verify a user's identity (authentication) on login, verify access privilege (authorization) on a per command basis, and record all users' activities in their logs (accounting). The log entries contain critical information which includes

- (i) the timestamp of the access request
- (ii) the IP address of the targeted network device (e.g., Loopback address)
- (iii) the IP address of the remote user requesting access
- (iv) the user's login ID
- (v) the command line executed
- (vi) other information such as user terminal, privilege level, and timezone.

To use a TACACS+ or RADIUS system in a network, all routers in the network need to be configured with the IP addresses of the servers – typically there are multiple replicated servers for redundancy. A large network can further be divided into multiple zones, for example, by the device type or by the autonomous system (AS) that they belong to. Different zones may contain different user account and privilege settings.

3 Characteristics of Normal Operation Activities

As described in the Introduction, our objective is to monitor all operational activities in the network and detect potential intrusions and misuses. We start by examining normal network operational activities recorded in the AAA logs. We focus on aspects that would best distinguish normal activities from actions that an external or internal attacker might take. In the following analyses, we use data collected from a global tier-1 ISP network which generates tens of millions of TACACS+ log entries from tens of thousands of routers per day.

3.1 Failed Login Attempts

The most intuitive way to separate potential attacks from legitimate accesses is to check whether they can readily pass authentication. Attackers may expose themselves by inputting wrong login credentials. However, it is also expected that legitimate users sometimes “fat finger” their login ID or password. Thus, we examine failed login attempts in normal AAA logs (using one month's data).

Figure 1 plots the cumulative distribution function (CDF) of the number of consecutive login attempts before a successfully authenticated login. We consider a login request within one minute of a preceding one with the same origin IP, the same login ID, and the same target networking device as a consecutive login. We observe that more than 99.992% of logins pass authentication the first time. More than 85% of the remaining ones input the correct credentials the next time, and it is extremely rare that a user fails more than five times before finally getting it right. The ratio of login failure is considerably lower than that typically seen in computer systems[18]. This is likely due to the predominance of logins generated by automated network management tools – a unique characteristic of network infrastructure operations. Figure 1 demonstrates the potential of alarming on intruders when a small number (e.g., 6) of repeated failed logins are observed.

The above type of monitoring can be defeated if the attacker has a list of valid login IDs and device names – they can use a different login ID or target a different device




Fig. 1. CDF of the number of consecutive login attempts before a success




Fig. 2. CDF of the number of consecutive failed login from a common IP

when an attempt fails. We can improve detection by looking for consecutive login failures from a common origin IP irrespectively of the login ID attempted. Figure 2 plots the CDF of the number of consecutive (i.e., less than one minute apart from the preceding one) login failures. We observe that around 85% of the rejected login attempts are either rectified or abandoned in six times or fewer. However, there are some login attempts going as high as a few hundreds in a row. Manual inspection finds that they were due to network management scripts running out of sync with router CLI (e.g. sending password when login ID is expected or *vice versa*). This rarely occurs, but when it does it produces many consecutive login failures – and should correctly trigger an alarm.

3.2 Login Access Pattern

As described in Section 2, an AAA log entry contains login access information characterized by the user login ID, the origin IP address, and the target router IP address. We define a *login session* as the network management activities sharing the common triple and being close in time (e.g., with an idle timeout of 10 minutes).

The login access information can be valuable in capturing attackers. For example, an origin IP that is not part of a block of addresses previously seen as an originating address in the logs is a strong indication that the network periphery protection may have a hole. Furthermore, each network operator typically has a rather stable set of work locations from which he/she manages the network, and due to his/her role, there can be a fixed set of network devices that the operator typically manages. So source and destination IP addresses will tend to be consistent over time for many operators.

We first look at the association of the operators’ login and the origin IP address. Figure 3 plots the CDF of the number of distinct origin IP addresses associated with a login ID in a month. We observe that 68% of login IDs manage the network from only one IP address. If we consider common origin subnets (with varying size), the number rises to 75% for /24 IP prefixes and 80% for /16 IP prefixes. In the rest of the paper, we will use /24 IP prefixes when aggregating origin IP address – it is not excessively large, yet can accommodate most of logins from the same facility/office.




Fig. 3. CDF of the number of origin IP (prefix) per login ID




Fig. 4. Stability of the login ID and origin IP prefix association

Figure 3 also shows that even with /24 origin IP prefixes, about 1% of the login IDs access the network from more than 10 distinct IP prefixes. Looking into those, we find that there are cases when an operator first logs on from a gateway server to a router, and then logs on to other routers from that router. The loopback IP address¹ of the first hop router appears as the origin IP for the second access session. While such “stepping-stone” access sessions are not common, they do occur – operators use this either for convenience or under certain network conditions, for example, when direct access to the other routers is unavailable. We can tighten our rule to deal with this situation by excluding sessions which originate on a router or switch. This removes “stepping stone” sessions from the analysis. The solid line in Figure 4 plots the number of distinct origin IP prefixes against the rank of login IDs – to protect proprietary information, we normalize the rank of login IDs to be between 0 and 1. We find above 78% of users have only one (non-stepping-stone) origin IP prefix and no one logs on from more than 4 distinct IP prefixes. This indicates that there exists a strong stability in the access pattern characterized by login ID and origin IP prefix combination – deviating from it can be a symptom of attacks. Figure 4 also plots the stability of this access pattern month by month – the shaded area indicates that the same login ID and IP prefix association has appeared in the preceding month. This demonstrates strong predictability based on past access behavior pattern (the unshaded area is mostly due to new users or infrequent users who only access the network in the second month).

Going back to the “stepping-stone” sessions, by matching the *ssh* command on the first hop router and the remote login request on the second router, we can reconstruct the chain of stepping-stones. Figure 5 plot the distribution function of the length of these chains and the outbound fan-out of these chains. It is evident that both attributes are bound by a small number (e.g., 7) in normal operational activities. In contrast, an intruder working from a compromised router may attempt to gain information from a large number of other routers, which is likely to produce long chains or high fan-outs. Watching those attributes closely can be an effective way to catch the intruder.

We next turn to the association between the networking device and the IP prefix from which management activities originate. The solid line in Figure 6 plots the number of

¹ Address assigned to a virtual interface commonly used for network management.




Fig. 5. CDF of the number of routers in a “stepping-stone” chain




Fig. 6. Stability of association between the routers and origin IP prefix




Fig. 7. Stability of the login ID and network zone association




Fig. 8. CDF of the maximum gap in days between consecutive logins per ID

distinct origin IP prefixes versus the rank of the networking device IDs. As in Figure 4, we normalize the ranks to be between 0 and 1 to protect proprietary information. We find that 48% of the routers are only controlled from hosts within one /24 IP block during a one-month period. These control activities are likely routine network auditing and health monitoring. A small portion of the network devices are managed from a small number of (e.g., 2-7) IP prefixes, which correspond to the network operation centers (NOCs) responsible for those devices. Furthermore, some cross-region access can be unexpected in normal operations – it is suspicious if an operator from a regional NOC in Japan requests access to a router serving IPTV in the USA. Catching abnormal associations between routers and origin IP prefixes can be an effective way of identifying such cases. The shaded area in Figure 6 shows the overlapping associations that have appeared in the preceding month; this demonstrates the predictability of these associations, as the overlapping is very significant.

Finally, we examine the association between login IDs and network devices. Many users or software tools have limited scope in terms of the networking devices managed. The solid line in Figure 7 plots the number of distinct network zones (described in Section 2) that each login ID has accessed in one month. We again normalize the x-axis to avoid disclosing the size of the operator work-force. We observe that the majority of login IDs have a very limited scope (e.g., less than 3 zones) while a few of high-tier

operators or software tool IDs access many zones. The stability of the login IDs' access pattern is depicted by the month over month comparison shown in the shaded area. We observe a strong predictability that can be utilized for detecting intrusions or misuses.

3.3 User Behavior

As mentioned in Section 2, different login IDs (corresponding to different operators or network management tools) have different roles/functions. Each user is likely to have a roughly stable behavior in access schedule (frequency), type of control (e.g., monitoring, or troubleshooting, or configuration change), and class of commands (e.g., SONET controller settings versus ACL configurations). Significant deviation from normal behavior can be a symptom of an account becoming compromised and an intruder impersonating the owner of the login account. In this subsection, we examine the properties of such user behaviors exhibited in normal network management activities.

We first examine the inter-session time distribution. Figure 8 plots the maximum difference in days between two consecutive logins from the same ID in a six-month period. We observe a wide variability among different login IDs. Many login IDs access the network on a regular basis, with a gap of at most a few days. But there are a considerable number of login IDs that only access the network occasionally. This suggests that it may be helpful to profile login IDs in different groups according to their access frequency.

Figure 9 shows the CDF of the average number of login sessions per login ID per day in a representative month. Here we exclude the days when the login ID is not active from the average statistic. The tail part of the curve, which goes several order of magnitude larger, is cut off so that we do not disclose the exact number of devices in the network. We observe that the majority of the login IDs have only a few login sessions per day. For example, 65% of IDs log onto the network no more than 5 times daily (on average). There are also many software tools and network management scripts producing over a hundred login sessions on daily basis. A login account suddenly changing its behavior, especially from having a small number of login sessions daily to a large number of them on a given day, is unusual or abnormal behavior and should be examined to see if it indicates a problem. Similarly, Figure 9 also shows the CDF of the average number of distinct networking devices accessed per login ID per day. Compared to number of sessions, it shows even more concentration – 65% of IDs log on to no more than 2 networking devices daily (on average). The tail portion of the curve again is dominated by software tools monitoring a large number of devices regularly, such as network configuration auditing tools. A surge in the number of distinct networking devices that a user initiates in a short period of time might be an intruder scouting for information. To understand the variability on this metric, Figure 10 shows a scatter plot of the coefficient of variation (CV) versus the mean – each point represents one login ID. We find that most of the CVs are bound by a small number (e.g., 3), while the login IDs with large number of average daily device accesses typically have much smaller CVs – suggesting that they can be more tightly bounded.

The set of router control commands and configurations used by a login ID is expected to exhibit some stability too. For a login account used by a software tool, the set of commands is determined by programming and rarely changes. For an operator's login,




Fig. 9. CDF of the mean number of sessions initiated per login ID per day




Fig. 10. CV versus mean of the number of distinct devices accessed per login ID per day




Fig. 11. Keywords Frequency Distribution




Fig. 12. Average Cosine similarity on the keywords frequency per login ID

the subset of commands should be subject to his/her privilege level and tightly related to his/her job role. However, extracting the exact association between login ID and the subset of the commands from knowledge about network operations is a challenging task – for example through code analysis of the software tools or meticulous review of all operation job functions.

In contrast, we take an approach that is based on historical data analysis and is detached from the semantics of router control commands, as follows: (a) we tokenize the commands (i.e., separate words in the command by white space); (b) we consider the tokens that contain any number as parameters (e.g., IP address) and remove those tokens; (c) we remove any non-alphabetic characters in each of the tokens and convert the remainder into lower case letters – we will refer to these as the *keywords*; (d) we profile each user with the set of keywords used.

Figure 11 shows the likelihood that a keyword is present in a command (sorted in decreasing order) based on one month of logs. We observe a strong skewness in the distribution, which can be well modeled by Zipf’s distribution. The high ranked keywords are those used in monitoring network health (e.g., ping, vrf, show). And most of the bottom ranked ones are some arbitrary tokens (such as customer name) referenced in the description field of certain router configurations or some misspellings (due to typos




Fig. 13. Detection System Architecture

by operators) of other keywords. It is sufficient to keep track a subset of keywords (e.g., top 1000) and represent the remainder simply as the *other* keyword. Figure 12 shows the stability of the use of keywords per login ID, which plots the average cosine similarity of the keyword frequency distribution comparing one day against the previous active day. Login IDs with a high number of daily log entries trend to have high predictability one day to another. Deviation from the regular command keywords, especially for a software tool account, can be a symptom of an intruder impersonating the owner of the account.

4 Design of an Online Intrusion Detection System

Based on the analysis from the previous section, we design an online intrusion detection system that oversees the network management activities of the ISP network and detects and alarms on anomalous patterns. Figure 13 shows the system architecture. We collect logs from AAA servers in near real-time. The logs are fed into an online preprocessing module, which extracts critical information and updates on an entry by entry basis the running states of sessions, login IDs, origins, and commands that are required for the different intrusion detection rules. Periodically (e.g., once a day) the running states are fed into an offline profiling module in which the different profiles required by the rules are updated – the initial profiles can be constructed via offline analysis of an extended period of historical data. The online rule checkers examine the running states against the profiles and rules and tag the corresponding log entries with a threat score. An aggregation module then sums the threat score in a window according to the login ID or origin IP. Finally, an alarming and presentation module makes the information available to network security operators.

4.1 Domain Knowledge-Based Rules

We first define a set of rules that is specific to the network under study. We maintain a list of the IP address blocks that belong to the ISP network and check the origin IP of each AAA log entry against the list. An IP address from outside of the network indicates a breach of the ISP’s periphery protection, and consequently the log entry is given a high threat score.

We also track the timestamp of the last login failure from each origin IP address and if a new failed login attempt is observed within T_1 seconds we update the timestamp

and increment the count of consecutive login failures for the origin IP. Once the count exceeds a threshold N_1 , we output the entries of these login attempts and assign a threat score to each of them. The timestamp and failure counts are reset when a successful login from the origin IP is made or the timeout T_1 is exceeded. With such a rule in place, an intruder that attempts to stay under the radar has to significantly slow down its attack, reducing the efficacy of the attack and prolonging the exposure.

There is also a rule in this category for stepping stone sessions. We trace stepping stone accesses as they occur and when the length of the access chain becomes greater than a threshold N_2 or the fan out becomes greater than N_3 , we assign a threat score to the sessions involved.

Finally, as a countermeasure to sophisticated intruders hiding themselves by disabling AAA logging right after gaining access to the network elements, we assign a high threat score to any log entry that modifies AAA logging settings (e.g., `tacacs-server` command on Cisco IOS). In the ISP network, changes to AAA logging settings should not happen in normal circumstances, so this rule should never generate a false trigger.

4.2 Rules Based on Access Pattern Profiles

In our daily association profile we keep track of the following attributes: (1) origin IP prefix (2) login ID (3) \langle login ID, origin IP prefix \rangle (4) \langle login ID, device zone \rangle (5) \langle origin IP prefix, device zone \rangle . For each entry we track the most recent date of appearance and the cumulative number of appearances. We delete an entry when the most recent appearance is more than T_2 (e.g., 180) days and add new entries to the long term profile once their count is sufficiently large.

We assign a threat score to sessions that do not match the existing profile. Note that if a session is from a new origin IP or new login ID, we do not include the threat score due to the lack of associations in (3), (4) and (5). The weight of the threat score of new associations of (3), (4) is set to be higher as the cumulative count of the login ID increases – our confidence to assert suspicious activities increases with more history data. Similarly, the weight of the threat score for (5) increases when the cumulative count for the corresponding origin IP prefix increases.

4.3 Rules Based on Statistical Models of the Access Profile

We track the mean and variance of the following attributes: (1) daily number of sessions per login ID, (2) daily number of distinct routers accessed per login ID, and (3) daily frequency count of command keywords per login ID for top N_4 and the *other* keywords.

We use the EWMA (Exponentially Weighted Moving Average) algorithm in estimating the running statistics for attribute X on day t :

$$Mean_t = \alpha X_t + (1 - \alpha) Mean_{t-1}$$

$$Var_t = \alpha (X_t - Mean_t)^2 + (1 - \alpha) Var_{t-1}$$

When computing the daily average, we exclude the case where the corresponding attribute is zero on day t – for example, when the user is inactive on the day.

If at any time, the daily cumulative counts reach or exceed a pre-calculated threshold for the attribute, we will assign a threat score to the access sessions involved.

The thresholds are determined as follows. Since the login IDs that have a high number of daily sessions (i.e., the highly active accounts) exhibit low variability as shown in Section 3.3, we set the thresholds in the same way as anomaly detection in Gaussian random variables: $\text{Threshold} = \text{Mean} + N_4 \times \text{Var}^{1/2}$. For the login IDs that have a moderate number of daily sessions, we set the threshold to be the product of a constant factor and the mean value: $\text{Threshold} = \text{Mean} \times N_5$. This is based on the observation that the coefficients of variation are bound by a small constant. And finally for the large portion of login IDs that only access the network occasionally, we set the threshold to be a small constant: $\text{Threshold} = N_6$. Once an attribute exceeds the defined threshold, a threat score is assigned. The value (weight) of the threat score is set according to a sublinear function of the corresponding attribute (the daily cumulative count) value.

4.4 Aggregation of Threat Scores

Using the rules above, in the course of a day, we maintain an updated set of AAA log entries that are assigned non-zero threat scores together with the rule triggered – to aid further examination by network security operators. Suspicious log entries can be noise (i.e., triggered by abnormal activities of low interest to security) – for example, an operator starting to use a new set of commands can trigger a violation detection by the user-keyword-rule. To reduce the chance that a network security operator has to investigate a non-critical violation, we further aggregate these log entries by login ID and by origin IP. The idea is that real attackers may be caught by multiple rules and by aggregating the threat scores on a per login ID or per origin IP basis they can be further distinguished from non-critical anomalies.

To achieve this, we use a moving window of T_3 (e.g., 1 day), and sum up the threat score within the window for all login IDs and origin IPs. We then set a threshold N_7 based on historical data. When we observe an aggregate threat score exceeding N_7 , we generate an alarm to the network security operators. We also display all suspicious activities on a dashboard report from which network security operators can pull information on demand.

5 Evaluation

We evaluate our system from two perspectives – the rate of anomalies detected from day-to-day network management activities and the effectiveness of detecting artificially injected anomalous activities. The former quantifies the resources required to investigate potential misuses and intrusions. The latter quantifies the chance that an anomaly goes undetected by our system.

5.1 Running System Performance

Figure 14 shows the distribution of the aggregate threat score in a month using two types of aggregations – by login ID (solid line) and by origin IP address (dashed line). We observe that about 93% of login IDs and 84% of origin IPs pass the system without raising




Fig. 14. CDF of aggregate threat score by login ID in one month




Fig. 15. CDF of aggregate threat score by login ID in one month w.r.t. different α values

any threat score. Meanwhile, there exist a small number of cases in which the system reports a high threat score. By manual inspection, we find most of them correspond to unusual network changes such as a newly deployed network management center or a major software upgrade on an existing network management tool. We will see an example of this in Section 5.3. We also find that a non-negligible fraction of cases, mostly with a relatively low threat score (e.g., less than 4), correspond to a small number of log entries from either a newly enabled or a very infrequently used account. Profiles for such accounts are difficult to construct based on history, and they typically do not generate many activities to drive the threat score high. We will propose and evaluate a system enhancement to further reduce the false alarms due to new user accounts in Section 6.

Figure 14 defines the tradeoff curves between the alarming rate and the sensitivity to anomalous activities. Raising the alarm threshold (the N_6 in Section 4.4), reduces the number of cases that security operators have to investigate but also reduces the chance of catching a stealthy intrusion. For a concrete example, setting N_6 to 5 would produce a few alarms per week on average, which is quite manageable for the network security operators.

We note that the above N_6 and several others as described in Section 4 are parameters used in the system. We do not present the exact values for the parameters in our running system due to security considerations. Instead, we show through an example our reasoning on parameter selection. Figure 15 shows the solid line in Figure 14 with a varying α value used in the EWMA estimate. Note the x -axis is in log scale. Different α values effectively factor in different amounts of history data. Setting $\alpha = 0.05$ effectively ignores (e.g., weight less than 0.01) data more than 90 active days old while $\alpha = 0.3$ effectively ignores data more than 13 active days old. However, Figure 15 shows there is little difference in the threat scores among the four different α values – indicating that a short history is sufficient for the system.

5.2 Controlled Experiment

Due to the lack of real attack observed in the network, we manually generate a set of data with anomalies for the performance test. Using Figure 14 as a reference point,

we design a controlled experiment as follows. We first randomly select 50 pairs of non-overlapping login IDs. We then take one day's worth of AAA logs from our running system and substitute the login ID field in the log entries such that the two login IDs in each of the chosen pairs are switched. Finally, we feed the manipulated AAA logs into our running system and monitor the output.

Figure 16 presents the CDF curves of the aggregate threat score based on the original AAA logs (dashed line) and the synthetic data (solid line) respectively. We find that our system is able to detect many of the behavior changes introduced by the login ID swapping. 72 out of the 100 login IDs report non-zero threat score and among them, 30 login IDs have a threat score higher than 10. This result highlights the sensitivity of our system – we do not expect one user impersonating another can always be caught (for example two operators with the same role may hardly be distinguishable), but compared to the baseline threat score distribution under normal operation, the synthetically generated ID-swappings stand out significantly.

For a closer look at how our system detects anomalous behaviors, Figure 17 and Figure 18 compare the contribution to the threat score from rules based on access pattern profiles with the contribution from rules based on statistical models of the access profile respectively. For each of the 100 login IDs, we plot their threat score against the difference in the daily average access frequency to the substituting login ID in the AAA logs. For example, a login ID *abc*, with 16 sessions per day on average, which is replaced by login id *xyz*, with an average of 1024 sessions per day, would have its threat score plotted at 6 (i.e. $\log(1024) - \log(16)$) on the *x*-axis. We observe that both the access pattern changes and the access statistics changes have contributed to the high threat score. The higher the difference in the amount of access activities between the pair of swapped login IDs, the higher the resulting threat scores – with the exception in Figure 18 on the negative side of the *x*-axis. The exception arises because the rules based on the statistical models are one-sided, i.e., we do not alert on a “busy” user suddenly becoming less active, as this behavior change does not seem to pose any security threat.

5.3 Operational Case Studies

We now look at an example in which our system alerted with a high threat score. Figure 19 plots the aggregate threat score of a particular login ID over the course of four days. The login ID is used by a software tool that periodically initiates `ping` commands among the various provider edge (PE) routers of the VPN customers to monitor their VPN health.

Starting in the afternoon of day 2 of the plot, we observe a fast increase in the threat score by the login ID. In less than two hours, the threat score passed the 99.5% alarming threshold and kept rising. It turned out that the software tool was upgraded that day and the new control sessions included a `show version` command that collects the router OS version across the network – similar to what might occur if an intruder attempted to collect information as preparation for attacks. After validating the change of behavior due to software upgrade, we included the pattern change in the profile update at the end of day, which greatly reduced the threat score on day 3. The corresponding statistical models were further updated at the end of day 3 and the new pattern then got fully captured by the profiles. Hence, there was no more threat score on day 4.




Fig. 16. CDF of aggregate threat score on real data and synthesized data




Fig. 17. Threat score by rules based on access pattern profiles




Fig. 18. Threat score by rules based on statistical models of access profile




Fig. 19. Threat score For Tool account A

We next take a look at an example where a system alert was triggered due to a new user appearing in the system. In this case, a new operator X started a session to a device between 1:30 and 1:59; since this was a newly created login ID, bearing no historical association to either the network zone or the origin IP, the session produced a threat score of 1.5 at the end of the session. Later on in the day, X logged on to four more devices in the same network zone between 3:20 and 3:40 in 21 different login sessions, pushing the threat score to 7.7 and triggering an alarm. A network security manager examined the alarm and checked X 's information from a corporate directory service. The threat was quickly dismissed when it was learned that X was a new member of the organization that manages the network zone in question. We will show next how we take the operational experience acquired in such scenarios and enhance our system specifically for newly provisioned login IDs.

6 A System Enhancement Handling New Login IDs

Operational experience has provided a valuable insight – when investigating a threat alert regarding a user login, network security officers often depend on the user's information in the corporate directory in addition to the access traces and profiles from our

system. Inspired by this, we look into how to utilize the corporate directory information to further reduce the false threat alerts.

6.1 Quantitative Analysis

We first create an automation capability that queries the corporate directory service for any given login ID. In particular, we obtain the corporate hierarchy information of the owner of the login ID by following the management chain up until reaching the president of the ISP. We then examine how users across the corporate hierarchy relate in terms of network access patterns and behaviors.

Figure 20 plots the CDF of the profile distance with respect to the access patterns (Section 4.2 (3), (4)) – measured by the l^2 norm of the difference in their profile vector between a pair of user login IDs. We define the corporate distance of a pair of users as the maximum number of hops for them to reach a common manager – two operators sharing the same direct supervisor would have a corporate distance of 1, and two sharing the same second-level-up manager would have a corporate distance of 2, etc. Different lines in Figure 20 correspond to user pairs of different corporate distances.

We observe that users that are organizationally close have similar access profiles. The further apart two users are in corporate distance, the more distinct their network access patterns are. When we focus on the case in which two operators are under the same direct manager (the solid line), we find that in 26% of the cases they have identical access pattern; and in 56% of the cases, their difference is no more than 3 (e.g., having access to 3 different network zones). This suggests that using the profiles of other team members can serve as a reasonable approximation when the operator’s own history is not fully established.

Figure 21 and Figure 22 plot the CDF of the absolute difference in the daily average number of sessions (Section 4.3 (1)) and the daily average number of distinct devices accessed (Section 4.3 (2)) respectively. In both cases, we observe decreasing trends in the behavioral similarity as users’ corporate distances increase, although the gaps are smaller compared to Figure 20. Focusing on the solid lines, we find that in 65% of the cases when users have a common direct reporting manager, their daily average number of sessions exercised differs by no more than 5 and their daily average number of distinct devices accessed differs by no more than 2. We have also observed a similar commonality in the profile of command token frequencies. These point to the good potential of bootstrapping profile building for newly provisioned login IDs using their peers’ profiles.

6.2 Profile Bootstrapping

As described in the operational case examples, network security managers sometimes dismiss system generated threats based on the additional information regarding the job function of the operator. Ideally, if domain knowledge of all different job functions and their expected access profiles were available, we could incorporate it into the system. However, such domain knowledge is implicit, highly distributed across various organizations, and evolving over time. Hence it is operationally challenging to acquire and




Fig. 20. CDF of access pattern profile distance for user pairs grouped by the corporate distance




Fig. 21. CDF of average difference in daily session count for user pairs grouped by the corporate distance




Fig. 22. CDF of average difference in daily distinct device count for user pairs grouped by the corporate distance

maintain. Instead, we propose to take a data-driven approach that does not require domain knowledge input, based on the analysis above.

Specifically, we organize user login IDs into classes based on the login owner’s management hierarchy. We then construct class-profiles based on the profiles of the members of the class. When a new login ID appears in the AAA log, the system queries the corporate directory service to obtain the user class to which the new login ID belongs. The class profile is then used as the bootstrapping profile for the user. The bootstrapping profile is replaced by the user’s own profile once adequate history has been observed for the login ID. This bootstrapping process aligns in principle with the model training techniques using scarce data [13].

To construct the class-profiles, we take the median statistic on each dimension (i.e., a feature in the profile of login ID) among all users (u) belonging to the class (C). Specifically,

$$\text{Profile}^C = \langle f_1^C, f_2^C, \dots \rangle, \text{ where } f_i^C = \text{median}_{u \in C} f_i^u$$

This is applicable to both binary flags (e.g., whether a login ID and network zone association exists) or numerical values (e.g., average number of sessions per day). Using the median is known to produce an estimator of cluster centroid that is robust to outliers[9].

6.3 Effectiveness of the Enhancement

Comparing the threat alert rate with and without profile bootstrapping, for the same month of data from Section 5.1, we find that profile bootstrapping has reduced the threat alerts (using a threshold $N_6 = 5$) by 60% for login IDs with inadequate history (less than or equal to seven active days of history). The overall threat alert rate (for all login IDs) is consequently reduced by 17% with profile bootstrapping.

7 Related Work

Our work falls into the area of IDS (Intrusion Detection Systems) in computer and networking security, which dates back to 1980 when Anderson[1] first proposed a

computer security surveillance system. Over time, the research area has become more active as the Internet grew in scale and application diversity and new security threats constantly emerged.

IDS broadly divide into two categories: host-based (HIDS)[7,8,4,10,15] and network-based (NIDS)[6,16,5,3] — HIDS typically rely on information about running processes to catch intrusions to computer host(s); NIDS typically analyze network traffic in order to detect attacks.

Another taxonomy of IDS is based on detection principles[17]: anomaly-based IDS (AIDS) [16,7,5,8,4,10,15] capture anomalous traffic or processes based on analysis of normal patterns. Signature-based IDS (SIDS) [11,14,3,6] use known signatures of attacks to alert on viral activities. Our work aligns with AIDS in principle.

Masquerader detection is a branch of IDS. A masquerader is an attacker who obtains a user's password, penetrates the access control system and impersonates a legitimate user. Lunt et al [7] designed IDES as the first IDS handling masquerader detection, using a simple yet effective statistical model. Recently, different machine learning techniques such as Genetic Algorithm [4], Naive Bayesian classification[10], and Support Vector Machine [15] have been applied in this area. In this study, we build user behavior models from access and command invocation patterns using statistical methods and alert based on deviation from the model. It remains as future work to evaluate whether more sophisticated machine learning algorithms can improve sensitivity and accuracy in our problem setting.

8 Conclusion

In this paper, we have studied the problem of protecting the networking infrastructure and the information available therein for large scale enterprise or ISP networks. We have proposed to enhance existing security measures with an intrusion detection system overseeing all network management activities. By analyzing device access logs collected via the AAA system in a global tier-1 ISP network, we have gained tremendous insights on the features that distinguish normal operational activities from rogue/anomalous ones. We have further developed a real-time intrusion detection system that builds statistical models to profile normal operational activities and alerts in real-time on any deviation from the profiles. Our evaluation demonstrates that this system effectively identifies potential intrusions and misuses with an acceptable overall alarm rate.

For future work, we would like to explore using more sophisticated machine learning techniques in addition to statistical methods to capture anomalous activities, using other network logs, and using other information such as network maintenance schedules to suppress alarms about “intended anomalies”. We are also interested in further introducing automated mitigation control based on detected anomalies to the AAA system such that an attack or intrusion can be stopped as early as possible.

Acknowledgments. This research is supported in part by the National Science Foundation under NSF grant CNS-0904901.

References

1. Anderson, J.P.: Computer security threat monitoring and surveillance. Technical Report James P Anderson Co Fort Washington Pa, p. 56 (1980)
2. Carrel, D., Grant, L.: The TACACS+ protocol (January 1997)
3. Dreger, H., Feldmann, A., Mai, M., Paxson, V., Sommer, R.: Dynamic application-layer protocol analysis for network intrusion detection. In: Proceedings of the 15th conference on USENIX Security Symposium, vol. 15. USENIX Association, Berkeley (2006)
4. Iglesias, J.A., Ledezma, A., Sanchis, A.: Creating User Profiles from a Command-Line Interface: A Statistical Approach. In: Houben, G.-J., McCalla, G., Pianesi, F., Zancanaro, M. (eds.) UMAP 2009. LNCS, vol. 5535, pp. 90–101. Springer, Heidelberg (2009)
5. Krishnamurthy, B., Sen, S., Zhang, Y., Chen, Y.: Sketch-based change detection: methods, evaluation, and applications. In: Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement, IMC 2003, pp. 234–247. ACM, New York (2003)
6. Li, Z., Xia, G., Gao, H., Tang, Y., Chen, Y., Liu, B., Jiang, J., Lv, Y.: Netshield: massive semantics-based vulnerability signature matching for high-speed networks. In: Proceedings of the ACM SIGCOMM 2010 Conference on SIGCOMM, SIGCOMM 2010, pp. 279–290. ACM, New York (2010)
7. Lunt, T.F., Jagannathan, R., Lee, R., Listgarten, S., Edwards, D.L., Neumann, P.G., Javitz, H.S., Valdes, A., Lunt, T.F., Jagannathan, R., Lee, R., Listgarten, S., Edwards, D.L., Neumann, P.G., Javitz, H.S., Valdes, A.: Ides: The enhanced prototype - a real-time intrusion-detection expert system. Tech. rep., SRI International, 333 Ravenswood Avenue, Menlo Park (1988)
8. Maggi, F., Matteucci, M., Zanero, S.: Detecting intrusions through system call sequence and argument analysis. *IEEE Transactions on Dependable and Secure Computing* 7, 381–395 (2010)
9. Maronna, R., Martin, R., Yohai, V.: Robust statistics: theory and methods. Wiley series in probability and statistics. J. Wiley (2006)
10. Maxion, R.: Masquerade detection using enriched command lines. In: Proc. of 2003 International Conference on Dependable Systems and Networks, pp. 5–14 (June 2003)
11. Paxson, V.: Bro: a system for detecting network intruders in real-time. *Comput. Netw.* 31, 2435–2463 (1999)
12. Rigney, C., Willens, S., Rubens, A., Simpson, W.: Remote authentication dial in user service, radius (2000)
13. Robertson, W., Maggi, F., Kruegel, C., Vigna, G.: Effective Anomaly Detection with Scarce Training Data. In: Proceedings of the Network and Distributed System Security Symposium (NDSS), San Diego, CA (February 2010)
14. Roesch, M.: Snort - lightweight intrusion detection for networks. In: Proceedings of the 13th USENIX Conference on System Administration, LISA 1999, pp. 229–238. USENIX Association, Berkeley (1999)
15. Salem, M.B., Stolfo, S.J.: A comparison of one-class bag-of-words user behavior modeling techniques for masquerade detection. *Security and Communication Networks* (2011)
16. Song, Y., Keromytis, A.D., Stolfo, S.J.: Spectrogram: A mixture-of-markov-chains model for anomaly detection in web traffic. In: NDSS. The Internet Society (2009)
17. Stefan, A.: Intrusion detection systems: A survey and taxonomy. Technical Report 99(Technical report 99-15), 1–15 (2000)
18. Suo, X., Zhu, Y., Owen, G.S.: Graphical passwords: A survey. In: Proceedings of the 21st Annual Computer Security Applications Conference, pp. 463–472. IEEE Computer Society, Washington, DC (2005)