Monitoring for Slow Suspicious Activities Using a Target Centric Approach

Harsha K. Kalutarage, Siraj A. Shaikh, Qin Zhou, and Anne E. James

Digital Security and Forensics (SaFe) Group, Faculty of Engineering and Computing
Coventry University, Coventry, CV1 5FB, UK
{kalutarh,aa8135,cex371,csx118}@coventry.ac.uk

Abstract. Slow, suspicious and increasingly sophisticated malicious activities on modern networks are incredibly hard to detect. Attacker tactics such as source collusion and source address spoofing are common. Effective attribution of attacks therefore is a real challenge. To address this we propose an approach to utilise destination information of activities together with a data fusion technique to combine the output of several information sources to a single profile score. The main contribution of the paper is proposing a radical shift to the focus of analysis. Experimental results offer a promise for target centric monitoring that does not have to rely on possible source aggregation.

1 Introduction

Slow, suspicious and increasingly sophisticated malicious activities on modern networks are incredibly hard to detect. An attacker may take days, weeks or months to complete an attack life cycle. Attacks may blend into the network noise in order to never exceed detection thresholds and to exhaust detection system state. A particular challenge is to monitor for such attempts deliberately designed to stay beneath detection thresholds. Attacker tactics such as source collusion and source address spoofing are common, and therefore make such attacker detection very hard. To address this we propose a method that does not require correlating to a common source. We shift the focus away from potential sources of attacks to potential targets of such activity. The proposed approach is designed to utilise destination information of activities together with a data fusion technique to combine the output of several information sources to a single profile score. We analyse for suspicious activities based on (or around) the destination information of the activities only.

2 Methodology

The problem of *target-centric* monitoring is broken down into two sub problems: *profiling* and *analysis*. Profiling provides for evidence fusion across spaces and accumulation across time, updating the normal node profiles dynamically based on changes in evidence. A multivariate version of simple Bayesian formula is

A. Bagchi and I. Ray (Eds.): ICISS 2013, LNCS 8303, pp. 163–168, 2013.

used as the method for evidence fusion to profile destination of activities during a smaller observation window. Those short period profiles are accumulated over the time to generate profiles for extended period of times (larger windows). It reduces the sheer volume of information such as raw logs and events [1], provided by number of different type of sources (e.g. SIDSs, anomaly detection components, file integrity checkers, AV, information from L3 switches), to a single value profile score for each node. Analysis distinguishes between anomalous and normal profiles using Grubbs' test [2]. Let H_1 and H_2 be two possible states of a node in computer network. We define H_1 as a node under attack and H_2 as a node not under attack. H_1 and H_2 are mutually exclusive and exhaustive states. P(H_1) expresses the belief, in term of probability, that the node is in state H_1 in absence of any other knowledge. Once we obtain more knowledge on our proposition H_1 through multiple information sources, in form of evidence E={e_1,e_2,e_3,...,e_m}, our belief is expressed as a conditional probability $p(H_1/E)$. Using Bayes theorem, and assuming statistical independence between information sources:

$$p(H_1/E) = \frac{\prod_{j=1}^{m} p(e_j/H_1).p(H_1)}{\sum_{i=1}^{2}\prod_{j=1}^{m} p(e_j/H_i).p(H_i)} \tag{1}$$

The assumption on statistical independence above is reasonable as we propose to use distinct types of information sources, which operate independently. In practice, a good Security Information Event Management (SIEM) deployment aggregates a number of solutions from many independent vendors [1]. When likelihoods $p(e_j/H_i)$ and priors $p(H_i)$ are known, the posterior $p(H_1/E)$ can be calculated. $p(H_1/E)$ terms in Equation 1 can be accumulated by time and used as a metric to distinguish targeted nodes from other nodes. We use the univariate version of Grubbs' test [2] to detect anomalies points in a given set of node profiles, subject to the assumption that *normal* node profiles in a given set follow an unknown Gaussian distribution. For each profile score x, its z score is computed as $z = \frac{|x-\bar{x}|}{s}$; where $\bar{x}$ and s are the mean and standard deviation of data set. A test instance is declared anomalous at significance level α if

$$z \geq \frac{N-1}{\sqrt{N}} \sqrt{\frac{t^2_{\alpha/N,N-2}}{N-2+t^2_{\alpha/N,N-2}}} \tag{2}$$

where N is the number of profiles points in the set, and $t_{\alpha/N,N-2}$ is the value taken by a t-distribution (one tailed test) at the significance level of $\frac{\alpha}{N}$. The α reflects the confidence associated with the threshold and indirectly controls the number of profiles declared as anomalous [3]. This is a vertical analysis to detect one's aberrant behaviour with respect to her peers.

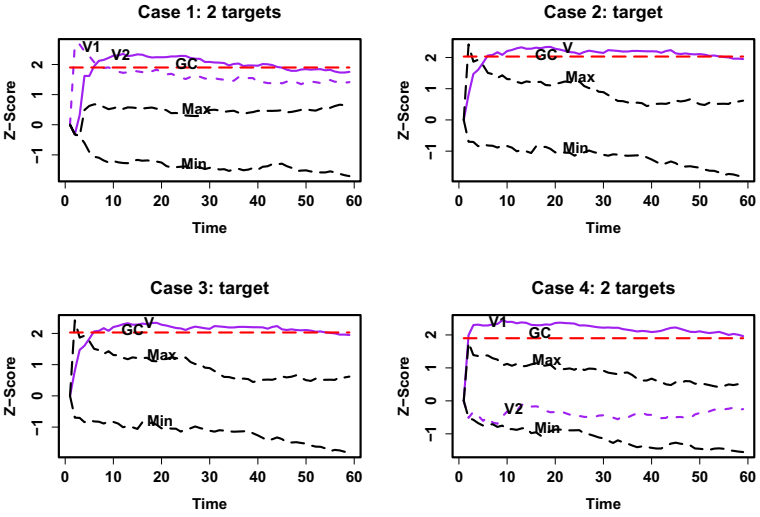

Fig. 1. Monitoring utilising the destination information of slow activities

3 Experimental Analysis

Network simulator $NS3$ is used to build a network topology consisting of a server farm and 10 subnets of varying size. Anonymous attackers, located in 3 subnets, are launching slow attacks on nodes in the server farm in a random manner. Anomalous traffic, by means of unusual port numbers, is generated along with normal traffic within and between networks using a Poisson arrival model. To simulate innocent events like user mistakes, suspicious traffic is also generated by normal nodes, but at different rates. If λ_a, λ_n are mean rates of generating suspicious events by attacker and innocent nodes respectively, we ensured maintaining $\lambda_a = \lambda_n \pm 3\sqrt{\lambda_n}$ and $\lambda_n (\leq 0.1)$ sufficiently smaller for all our experiments to characterise slow suspicious activities which aim at staying beneath the threshold of detection and hiding within the background noise. The idea to use the above relationship for generating attacker activities was to keep them within the *normalilty range* of innocent activities (i.e. background noise). $\sqrt{\lambda_n}$ is the standard deviation of rates of suspicious events generated by normal nodes. Each simulation was run for a reasonable period of time to ensure that enough traffic is generated. For the purposes of simulation prior probabilities $p(H_1) = \frac{1}{2}$ and likelihoods $p(e_j/H_1) = k$ (arbitrary values $\geq$0.5 and $\leq$1) were used to distinguish different types of events. Estimation of these probabilities for real networks can be found in [3,4,5].

Figure 1 shows how the targets of slow attacks are detected. In case 1 two nodes on the server farm are targeted by three attackers, in case 2 one node on the server farm is targeted by three attackers, in case 3 one node on the server farm is targeted by a single attacker, and in case 4 two nodes on the server

farm are targeted by one attacker. The Bayesian model from equation 1 is used to generate profile scores for obtaining results in Figures 1 and 2. Results in Figure 1 are obtained utilising target information while results in Figure 2 are obtained utilising source information. In Figure 2, similar results were obtained for all three attackers in cases 1 and 2, but the results for only one attacker are presented due to space constraints. The same trace is used for obtaining results in both Figures. *Min*, *Max* and *GC* are the minimum, maximum and Grubbs' critical value (i.e. the threshold) for profile scores of normal nodes in each subnet where a targeted node (V in Figure 1) and an attacker node (A in Figure 2) are located.

Our approach is capable of detecting targets under attack successfully (see Figure 1). Targets cut off (or very close to) the threshold while normal nodes in target's subnet are significantly away from the threshold. As depicted in Figure 2, attackers hide among normal nodes, and the source-centric approach fails to detect them as quickly. Case 4, where colluded activities are not simulated, is an exception here as it detects only one target out of two. But in case 1, both target nodes are detected; a minimum number of observations are required in order to detect a target successfully. In case 1, since three attackers target two victims, there is a better chance for the monitoring system to observe enough evidence against each victim than it is in case 4. Finding the relationship between detectability and minimum number of observations required is future work.

Detection potential measures how likely an activity could be detected as a suspicious slow activity. It is expressed in terms of deviations of profile scores from the threshold line. The higher the deviation the better the chances of detection. On that basis the detection potential d is defined as: $d = z - GC$. Figure 3 compares this across the two approaches in each case. A (or A_i) represents the detection potential for attackers while V (or V_i) represents the detection potential for targets. The latter has a higher detection potential in all cases. Higher variations (fluctuations) on detection potential indicate a higher chance for false alarms.

4 Related Work

[6] offers a different direction for security monitoring by proposing a class of scanning detection algorithms that focus on what is being scanned for instead of who is scanning. But such an approach is not completely independent from the source information either. It uses the source information of scan packets for victim detection. Our approach does not require any information about the source. It completely depends on destination information and allows for any suspicious event on the network to be accounted for. Most importantly, we acknowledge two types of uncertainties of events defined in [7,4] in a Bayesian framework. Hence our effort is completely different from [6], but has been inspired from that work. Using Bayesian technique and its variants for intrusions detection can be found in [5]. The relevance of information fusion for network security monitoring has been widely discussed [3,8]. Our method is based on anomaly detection.

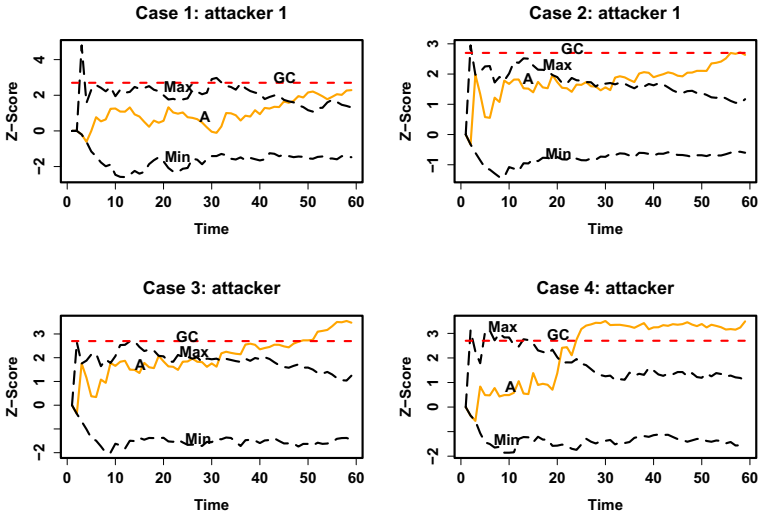

Fig. 2. Monitoring utilising the source information of slow activities

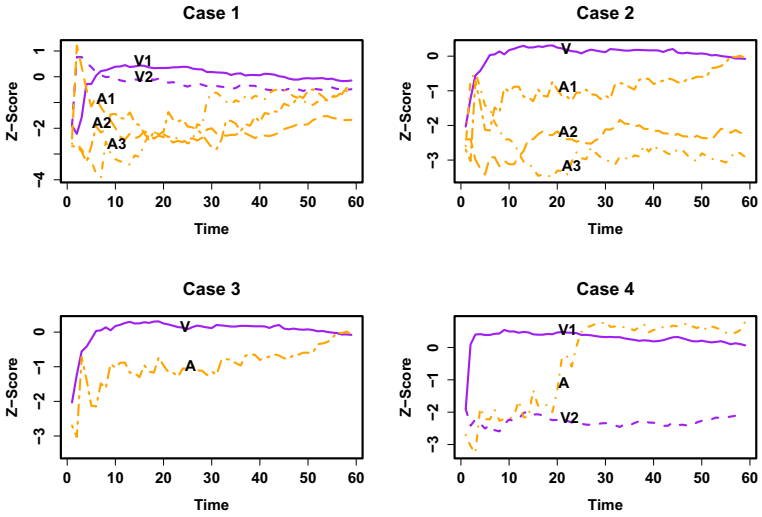

Fig. 3. A comparison of detection potential for each case

A number of other anomaly based detection approaches have been proposed, but most are general in nature [9,10,11]. Most of current incremental anomaly detection approaches focus on rapid attacks, have high rate of false alarms, are non-scalable, and are not fit for deployment in high-speed networks (refer to survey paper [12]), whereas our focus is on slow attacks.

5 Conclusion

One difficulty with attribution discussed earlier is that attacks are carried out in multiple stages using compromised machines as stepping stones (or in the form of bot-nets). The focus on targeted nodes takes into account the importance of preventing such compromise, which in itself should help to undermine attacks. One argues that monitoring systems could be deployed to achieve both attribution and early warning for attacks on target nodes. While this is feasible in theory, in practice this means the cost of monitoring is incredibly high, as networks expand in size, traffic volume rise, and slow attackers get slower. The main contribution of the paper is proposing a radical shift to the focus of analysis. We utilise a data fusion algorithm to combine the output of several information sources to a single score. It acts as a data reduction method and enables us to propose a lightweight monitoring scheme for the problem which is essential in near-real-time analysis of slow, sophisticated targeted attacks.

References

1. (CSIEM): Cisco security information event management deployment guide (August 2013), http://www.cisco.com
2. Grubbs, R.E.: Procedures for Detecting Outlying Observations in Samples. Technometrics 11(1), 1–21 (1969)
3. Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: A survey. ACM Comput. Surv. 41(3), 15:1–15:58 (2009)
4. Kalutarage, H.K., Shaikh, S.A., Zhou, Q., James, A.E.: Sensing for suspicion at scale: A bayesian approach for cyber conflict attribution and reasoning. In: 4th International Conference on Cyber Conflict (CYCON), pp. 1–19 (2012)
5. Siaterlis, C., Maglaris, B.: Towards multisensor data fusion for dos detection. In: ACM Symposium on Applied Computing, pp. 439–446 (2004)
6. Whyte, D., van Oorschot, P.C., Kranakis, E.: Exposure maps: removing reliance on attribution during scan detection. In: Proceedings of the 1st USENIX Workshop on Hot Topics in Security, HOTSEC 2006. USENIX Association (2006)
7. Kalutarage, H.K., Shaikh, S.A., Zhou, Q., James, A.E.: Tracing sources of anonymous slow suspicious activities. In: Lopez, J., Huang, X., Sandhu, R. (eds.) NSS 2013. LNCS, vol. 7873, pp. 122–134. Springer, Heidelberg (2013)
8. Vokorokos, L., Chovanec, M., Látka, O., Kleinova, A.: Security of distributed intrusion detection system based on multisensor fusion. In: 6th International Symposium on Applied Machine Intelligence and Informatics, pp. 19–24 (2008)
9. Patcha, A., Park, J.M.: An overview of anomaly detection techniques: Existing solutions and latest technological trends. In: Computer Networks. Elsevier (2007)
10. Kumar, S., Spafford, E.H.: An application of pattern matching in intrusion detection. In: Technical Report CSDTR-94-013 Purdue University, IN, USA (1994)
11. Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: A survey. ACM Computing Surveys 41 (2009)
12. Bhuyan, M.H., Bhattacharyya, D., Kalita, J.K.: Survey on incremental approaches for network anomaly detection. International Journal of Communication Networks and Information Security 3(3), 226–239 (2012)