How do we effectively monitor for slow suspicious activities?

Harsha K. Kalutarage, Siraj A. Shaikh, Qin Zhou and Anne E. James
{kalutarh, aa8135, cex371, csx118}@coventry.ac.uk

Digital Security and Forensics (SaFe) Research Group
Department of Computing, Faculty of Engineering and Computing
Coventry University
Coventry, CV1 5FB, UK

Abstract As computer networks scale up in size and traffic volume, detecting slow suspicious activity, deliberately designed to stay beneath the threshold, becomes ever more difficult. Simply storing all packet captures for analysis is not feasible due to computational constraints. Detecting such activity depends on maintaining traffic history over extended periods of time, and using it to distinguish between suspicious and innocent nodes. The doctoral work presented here aims to adopt a Bayesian approach to address this problem, and to examine the effectiveness of such an approach under different network conditions: multiple attackers, traffic volume, subnet configuration and traffic sampling. We provide a theoretical account of our approach and very early experimental results.

1 Introduction

In the domain of computer security, an attacker may take days, weeks or months to complete the attack life cycle against a target host. This allows for such activity to blend into the network as noise. There is no clear definition to distinguish slow attacks from typical attacks. Of interest to us is any suspicious attempt to stay beneath intrusion detection thresholds. Detection of such 'low and slow' attacks pose a number of challenges. Simply storing and analysing all full content capture is not feasible due to computational constraints. Research addressing this area uses incremental anomaly detection approaches. Most of current incremental anomaly detection approaches have high rate of false alarm, are non-scalable, and are not fit for deployment in high-speed networks [MBK11]. Chivers et al. use Bayes formula to identify slow insider activities [CNS+09,CNS+10]; Phillip et al.'s work is similar [BBSP04]. In [BBSP04] user behaviour is profiled and used to identify those who require further investigations. Kalutarage et al. propose an aproach for cyber conflict attribution and reasoning in a Bayesian framework, together with concept of statistical normality [KSZJ12]. Streilein et al. use multiple neural network classifiers to detect stealthy probes [SCW02]. Evidence accumulation as a means of detecting slow activities is proposed in [T.H02]. Schultz et al. claim that profiling suspected insiders provides one of the best ways of reverse engineering an attacker [ER01].

2 How do we effectively monitor for slow suspicious activities?

Our aim is to study effective monitoring for slow suspicious activity that can be investigated further. We break our research objectives down to the following structure.

1. How do we effectively attribute slow suspicious activity to the source?
   The main goal here is to establish means of attributing such activity. At this stage we are interested in determining different possible methods to achieve this, and particularly investigating whether a Bayesian approach is feasible.

2. What effect does varying network parameters have over effective monitoring?

[2] In: A. Editor, B. Coeditor (eds.): Proceedings of the XYZ Workshop, Location, Country, DD-MMM-YYYY, published at http://ceur-ws.org

We are particularly interested in studying subnet size and traffic volume, and how that may effect our ability to distinguish such activity. We will draw from this network design principles for more effective monitoring.

3. How do we effectively detect the target of such activity?
We acknowledge that the use of botnets and distributed sources makes it very difficult to attribute attacks. Of further interest is to determine the target of such activity. We will investigate methods to profile such nodes. Such methods need to be effective for scalable networks.

4. What effect does using sampling techniques has as a logging method?
Traffic volumes will continue to increase. This makes it ever more difficult to process and effectively monitor slow activity. Since we are not detecting for strict traffic signatures, we wish to investigate traffic sampling methods and evaluate their suitability for security monitoring of slow attacks.

3 Research Methodology

We look at the problem as two sub problems: *profiling* and *analysis*. Profiling is the method for evidence fusion across space and accumulation across time, which updates the normal node profiles dynamically based on changes in evidence. Analysis is the method for distinguishing between anomalous and normal profiles using statistical normality. We propose to use elements of network flow data as input to our profiling method. Flow data contains network and port addresses, protocols, date and time, and amount of data exchanged during a session. We use a multivariate approach to analyse such records. So for example suspicious port scanning activity may have the following characteristics: a single source address, one or more destination addresses, and target port numbers increasing incrementally. When fingerprinting such traffic, we examine multiple elements and develop a hypothesis for the cause of behaviour on that basis. We use a Bayesion approach to achieve this.

3.1 Building the hypothesis

The posterior probability of the hypothesis H_k given that E, is given by the well known Bayes' formula:

$$p(H_k/E) = \frac{p\left(E/H_k\right).p(H_k)}{p(E)} \tag{1}$$

Let H_k : hypothesis that kth node is an attacker, E_i is a flow record element and $E = \{E_1$=e$_1$, E_2=e$_2$, E_3=e$_3$,...,E_m=e$_m\}$ is the set of all suspicious evidence observed against node k during time t from m different independent observation spaces. Here $P(E)$ is the probability of producing suspicious events by node k, but on its own is difficult to calculate. This can be avoided by using the law of total probability. For independent observations, the joint posterior probability distribution can be obtained from (1) as:

$$p(H_k/E) = \frac{\prod\limits_{j} p(e_j/H_i).p(H_k)}{\sum\limits_{i}\prod\limits_{j} p(e_j/H_i).p(H_i)} \tag{2}$$

To calculate the posterior probability of node k being an attacker $p(H_k/E)$, it is necessary to estimate:

1. the likelihood of the event E given the hypothesis H_i, $p(E/H_i)$ and,
2. the prior probability $p(H_i)$, where $n \geq i > 0$.

Assuming that prior and likelihoods are known, (2) facilitates to combine evidence from multiple sources (all E_is) to a single value (posterior probability) which describes our belief, during a short observation period, that node k is an attacker given E. Aggregating short period estimations over time helps to accumulate relatively weak evidence for long periods. This accumulated probability term, $\sum\limits_{t} p(H_k/E)$ (t is time) known as profile value hereafter, can be used as a measurement of the level of suspicion for node k at any given time. These scores are converted into Z-scores for analysis.

A series of experiments have been conducted in a simulated environment to test the proposed approach. We use NS3 [NS311] to simulate our network and generate traffic patterns of interest, assuming a poison arrival model. Each simulation is run for a reasonable period of time to ensure that enough traffic is generated (over one million events). If λ_s, λ_n are mean rates of generating suspicious events (where we only generate a subset of flow data elements including source and destination address and port numbers, and where suspicious activity is judged by unexpected port numbers) by suspicion and normal nodes respectively, we ensure maintaining $\lambda_s = (\lambda_n \pm 3\sqrt{\lambda_n})$ 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 behind the background noise.

3.2 Early Results

Early results of our work are promising: our approach is able to distinguish multiple suspicious nodes from a given set of network nodes as shown in Figure 1.

We model *detection potential* D as a function of subnet size S and traffic volume V, where $D = k.(\frac{V}{b^S})^{\frac{1}{2}}$, and where k is a constant, which demonstrates the effect of varying the subnet size over ability to detect effective monitoring. This effect is demonstrated in Figure 2. The effects of total traffic volume on detection potential are also demonstrated in Figure 3. Relevant details for these results could be found in [KSZJ12].

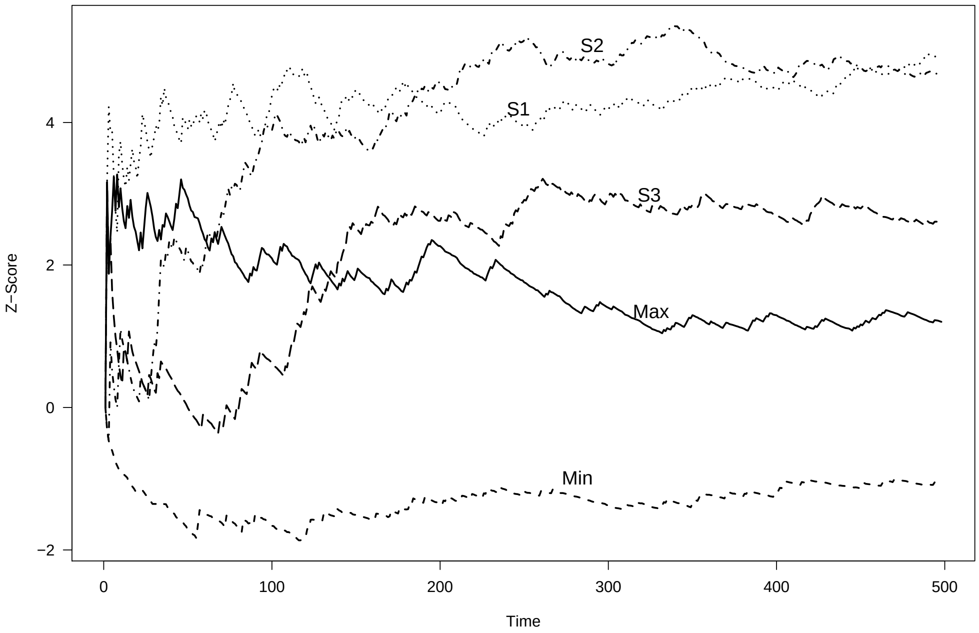

Figure 1: A Z-Score graph - for slow attack monitoring. S_i represents suspicious nodes. Min and Max represent the minimum and maximum Z-scores of normal nodes.

Our work aims to address the stated research goals by demonstrating how effective monitoring could be deployed in more realistic network topologies. We plan to continue with our experimental approach, and consolidate results towards the end to ensure a coherent and consistent picture emerges that is of practical value.

4 Contribution

This research aims to address a difficult problem. Monitoring infrastructures are overloaded both with data and tools. The question is: what do we with it? The difficulty is due to the increasing scale of networks, the diversity of user access provision to systems, the nature of suspicious activity and the corresponding need to monitor for serious attacks, and ultimately being able to effectively manage detection of intrusions.

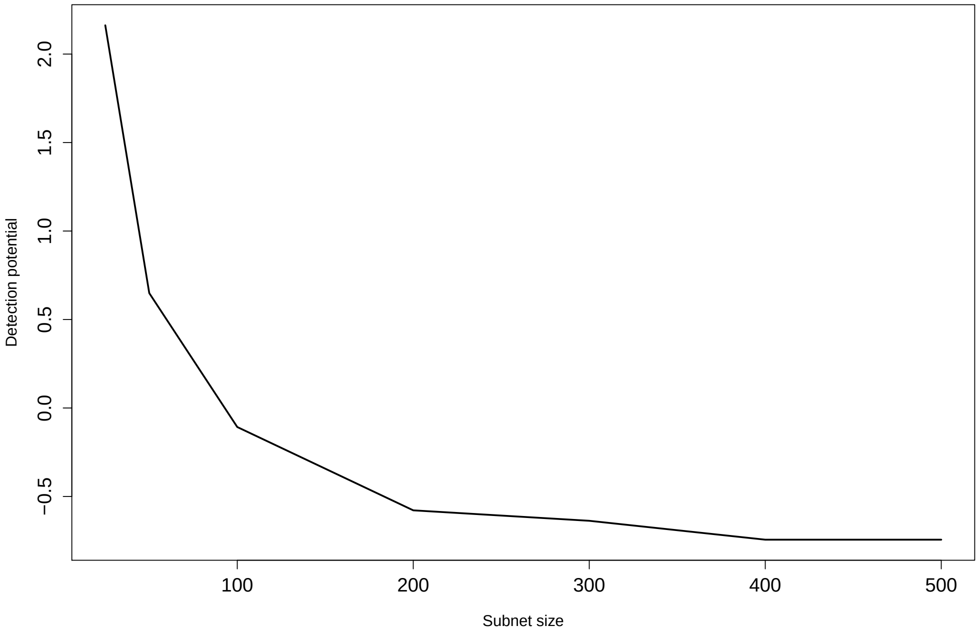

Figure 2: Smaller the subnet size, the better for detection⇒D $\propto \frac{1}{b^S}$, b is a constant.

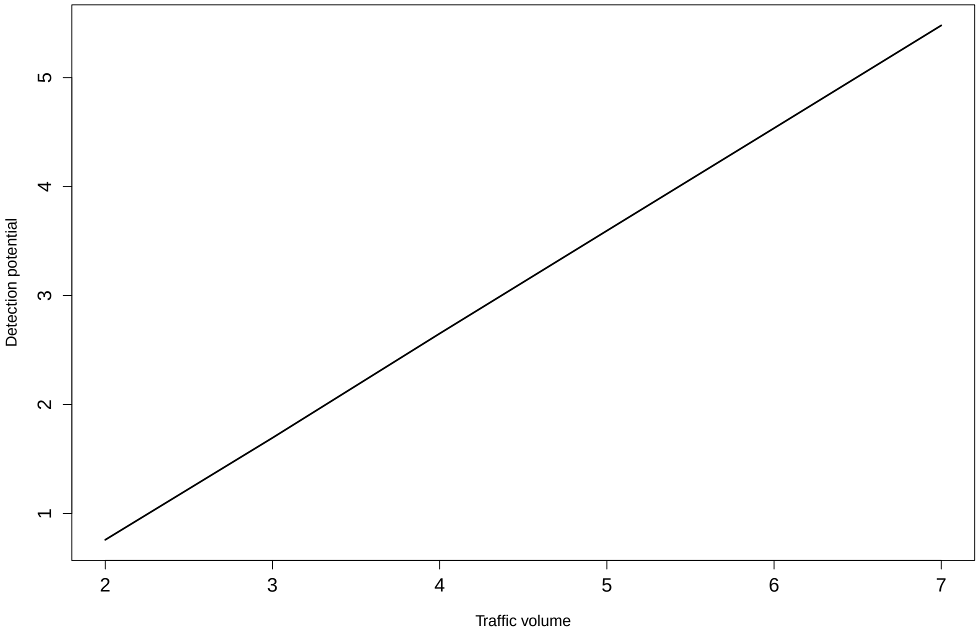

Figure 3: Higher the traffic volume, the better for detection⇒D $\propto V$.

Our ultimate goal is to offer a set of design principles and heuristics allowing for effective collection and analysis of data on networks. The first two research questions from Section 2 allow us to build defensible networks, where any source of suspicious activity could be detected effectively and quickly. This is about both better data analysis and network design. The third research question is inspired by related work investigating exposure maps [DOE06] and darkports [WvOK07], where we adapt our algorithm to profile target nodes for possible slow and suspicious activity. The underlying principle remains the same: we trade in state for computation. Ever increasing processing capacity increasingly makes this feasible. But traffic volumes indeed also pose a big challenge, and hence our final question is an attempt assess the feasibility of sampling traffic for analysis. This is also evidenced as feasible by some other work [BR12,PRTV10], and we propose to build on it.

Our aim is to remain domain agnostic. This allows for research to be applied at various levels, including better detection software, monitoring tools, and network design and configuration management solutions.

References

[BBSP04] Phillip G. Bradford, Marcus Brown, Bonnie Self, and Josh Perdue. Towards proactive computer-system forensics. In *In International conference on information technology: Coding and computing,IEEE Computer Society*, 2004.

[BR12] Karel Bartos and Martin Rehak. Towards efficient flow sampling technique for anomaly detection. In *Proceedings of the 4th international conference on Traffic Monitoring and Analysis*, TMA'12, pages 93–106, Berlin, Heidelberg, 2012. Springer-Verlag.

[CNS$^+$09] Howard Chivers, Philip Nobles, Siraj Ahmed Shaikh, John Clark, and Hao Chen. Accumulating evidence of insider attacks. In *(MIST 2009) (In conjunction with IFIPTM 2009) CEUR Workshop Proceedings*, 2009.

[CNS$^+$10] Howard Chivers, Philip Nobles, Siraj Ahmed Shaikh, John Clark, and Hao Chen. Knowing who to watch: Identifying attackers whose actions are hidden within false alarms and background noise. *Information Systems Frontiers, Springer*, 2010.

[DOE06] Whyte David, P.C.van Oorschot, and Kranakis Evangelos. Exposure maps: removing reliance on attribution during scan detection. In *Proceedings of the 1st USENIX Workshop on Hot Topics in Security*, pages 9–9, Berkeley, CA, USA, 2006. USENIX Association.

[ER01] E.E.Schultz and R.Shumway. Incident response: A strategic guide for system and network security breaches indianapolis. In *New Riders*, 2001.

[KSZJ12] Harsha K. Kalutarage, Siraj A. Shaikh, Qin Zhou, and Anne E. James. Sensing for suspicion at scale: A bayesian approach for cyber conflict attribution and reasoning. In *InProceedings of 4th International Conference on Cyber Conflict, NATO CCD COE*. NATO CCD COE Publications, Tallinn, June 2012.

[MBK11] M.H.Bhuyan, DK Bhattacharyya, and JK Kalita. Survey on Incremental Approaches for Network Anomaly Detection. *International Journal of Communication Networks and Information Security (IJCNIS)*, 3, 2011.

[NS311] NS3 Development Team. Ns3 discrete-event network simulator for internet systems, 2011.

[PRTV10] Antonio Pescap, Dario Rossi, Davide Tammaro, and Silvio Valenti. On the impact of sampling on traffic monitoring and analysis . In *Proceedings of 22nd International Teletraffic Congress (ITC) 2010*, pages 1–8, 2010.

[SCW02] William W. Streilein, Robert K. Cunningham, and Seth E. Webster. Improved detection of low-profile probe and novel denial-of-service attacks. In *Workshop on Statistical and Machine Learning Techniques in Computer Intrusion Detection*, 2002.

[T.H02] T.Heberlein. Tactical operations and strategic intelligence: Sensor purpose and placement. Technical Report TR-2002-04.02, Net Squared Inc, 2002.

[WvOK07] David Whyte, Paul C. van Oorschot, and Evangelos Kranakis. Tracking Darkports for Network Defense . In *Proceedings of Computer Security Applications Conference, 2007. ACSAC 2007.*, pages 161 – 171, 2007.