Early Warning Systems for Cyber Defence

Harsha Kalutarage, Siraj Shaikh, Bu-Sung Lee, Chonho Lee, Yeo Kiat

Early Warning Systems for Cyber Defence

Harsha Kalutarage[1], Siraj Shaikh[2], Bu-Sung Lee[3], Chonho Lee[3], and Yeo Chai Kiat[3]

[1] The Centre for Secure Information Technologies, Queen's University of Belfast, UK
h.kalutarage@qub.ac.uk
[2] The Centre for Mobility & Transport, Coventry University, UK
s.shaikh@coventry.ac.uk
[3] Nanyang Technological University, Singapore
{ebslee,leechonho,asckyeo}@ntu.edu.sg

Abstract Cybercriminals ramp up their efforts with sophisticated techniques while defenders gradually update their typical security measures. Attackers often have a long-term interest in their targets. Due to a number of factors such as scale, architecture and nonproductive traffic however it makes difficult to detect them using typical intrusion detection techniques. Cyber early warning systems (CEWS) aim at alerting such attempts in their nascent stages using preliminary indicators. Design and implementation of such systems involves numerous research challenges such as generic set of indicators, intelligence gathering, uncertainty reasoning and information fusion. This paper discusses such challenges and presents the reader with compelling motivation. A carefully deployed empirical analysis using a real world attack scenario and a real network traffic capture is also presented.

Keywords: Bayesian Inference, Cyber Defence, Cyber Warfare, Future Internet, Early Warning Systems

1 Introduction

Early warning systems for cyber defence is an emerging area of research which aims at alerting an attack attempt in its nascent stages. Wider definitions for a CEWS can be summarised as "a CEWS aims at detecting *unclassified* but potentially harmful system behaviour based on *preliminary indications* before possible damage occurs, and to *contribute* to an integrated and aggregated situation report" [1]. Although there can be many overlaps between a typical intrusion detection system (IDS) and a CEWS, a particular emphasis for a CEWS is to establish hypotheses and predictions as well as to generate advice on not yet understood (unclassified) situations based on preliminary indications [1]. In contrary, a typical IDS attempts to detect attack using *known indications of attack patterns* (these can be either signatures or anomalies) instead of using generic preliminary indications.

This paper identifies some research challenges compounded by the nature of computing for design and implementation of effective CEWSs, and discusses

potential solutions to overcome them. The paper starts with an empirical analysis in section 2. Section 3 presents research challenges. Related work is presented in section 4. Finally, section 5 includes a discussion with necessary suggestions to move forward in this research line.

2 An empirical analysis

The sole purpose of this section is to demonstrate the feasibility of using preliminary indicators to produce early warnings in a situation when the known indicators of attack pattern is not available to use typical IDS techniques. This proof of principle case study built on an existing work [2].

2.1 Attack scenario

We analyse the data set described in section 2.2 for heartbleed exploits attempts. The heartbleed vulnerability [3] lies in the implementation of heartbeat protocol extension of the transport layer security (TLS). Heartbeat consists of two message types: heartbeat request and heartbeat response. When a request message received, the receiver must send a corresponding response message carrying an exact copy of the payload of the request by allocating a memory buffer as:

```
buffer = OPENSSL_malloc(1 + 2 + payload + padding)
```

There was no length check for this memory allocation in OpenSSL 1.0.1 and prior. Hence an attacker can specify higher payload values than the actual payload in the request and hence abuse the server to read arbitrary memory locations. This allows attackers to read sensitive memory (e.g. cryptographic keys and credentials) from vulnerable servers. Since there is a maximum boundary for the total length of a heartbeat message, in a heartbleed attack attempt, a higher number of message frequency can be expected during a connection in order to leak as much as possible data from the server's memory. It should be noted that it is necessary to look at and compare two fields k_1 and k_2 (see Table 1) in the TLS layer data to detect exploits attempts. If $k_1 > k_2$ then it is explicitly a heartbleed packet.

2.2 MAWI data set

In the MAWI data set, traffic traces have been collected during a 15 minutes period on each day at several sampling points within WIDE backbone - an operational testbed network in Japan [4]. After removing privacy information, traces have made open to the public. Hence traces consist of only protocol headers. Readers are invited to notice the limitations of details in the MAWI dataset with respect to heartbleed detection. Since TLS layer data is not available in the data set, it is not possible to explicitly check for heartbleed attack atempts. Therefore our aim is to analyse the dataset for the same attack scenario using a set of preliminary indicators as defined in Table 1.

Known indicators to explicitly detect heartbleed exploit attempts by a typical IDS	**Preliminary indicators to implicitly warn heartbleed exploit attempts by a CEWS**
k_1 - requested payload length	i_1 - number of TCP segments from client to server
	i_2 - upload during a session
k_2 - actual payload length	i_3 - downoad during a session
	i_4 - time gap between two consecutive packets

Table 1: Two kind of indicators defined over a 443 session.

2.3 Mathematical basis

A node score is computed as follows. Let H be the hypothesis that given node (or IP address) is a heartbleed attacker and $I = \{i_1, i_2, i_3, i_4\}$ is a set of indicators defined within a 443 session (see Table 1). Assuming statistical independence between indicators and using well known log likelihood ratio,

$$ln\frac{P(H/I)}{P(\neg H/I)} = ln\frac{P(H)}{P(\neg H)} + \sum_k ln\frac{P(i_k/H)}{P(i_k/\neg H)} \tag{1}$$

During a smaller time window w, if $ln\frac{P(H/I)}{P(\neg H/I)} > 0$ then H is accepted. The rationale behind variable selection is that they are weakly connected to the behaviour of the heartbeat protocol. Our idea is to compare probability distributions of these variables in two populations, i.e. attack and clean, using equation 1. The prior belief $P(H)$ and $P(\neg H)$ were defined as follows.

$$p(H) = \begin{cases} 0.6, \text{ if the target node was doing a scan (port/host) prior to a session} \\ 0.4, \qquad\qquad\qquad\qquad\qquad \text{otherwise} \end{cases} \tag{2}$$

The likelihood distributions $P(i_k/H)$ and $P(i_k/\neg H)$ were estimated using "malicious" and "clean" datasets respectively. Distribution for each variable were proposed by looking at their histograms. A dataset prior to December 2011 (i.e. before the bug was introduced in Open SSL) was chosen as the "clean" set. The "malicious" set was chosen based on our assumption that there is a higher chance for heartbleed attack attempt during the heartbleed public announcement period. This is due to practical constrains accessing for a sufficiently large known heartbleed dataset.

2.4 Experimental setup & outcomes

Fifteen minutes duration traces from each day was split into 90 segments, each segment is a 10 second window. Within a window, nodes are profiled using equation 1. If a node obtained negative (-) scores throughout the observation period then that node is defined as "innocent". If a node obtained at least one positive (+) score during the observation period then that node is defined as a "suspicious" node. Among suspicious nodes, if any node stands out from their

peer nodes (i.e. beyond three z-scores) then that node is identified as "most suspicious". Zero (0) means the target node has not produced any event that are of interest to this analysis during the observation window.

Figure 1 presents experimental outcome during heartbleed public announcement period. The graph presents the node score against the time line. Note that 11 and 322 nodes (out of 9087 nodes) were selected as most suspicious and suspicious nodes respectively in which proposed analysis has reduced the search space by 96%. In order to understand the recurrent of the target scenario by the same or different nodes, above analysis is repeated periodically (every two months) since July 2011 to July 2014. Due to the space constraint, only two graphs at the beginning and two graphs at the end of the analysis are presented in figure 2. A detailed description of the analysis can be found in our earlier paper [5].

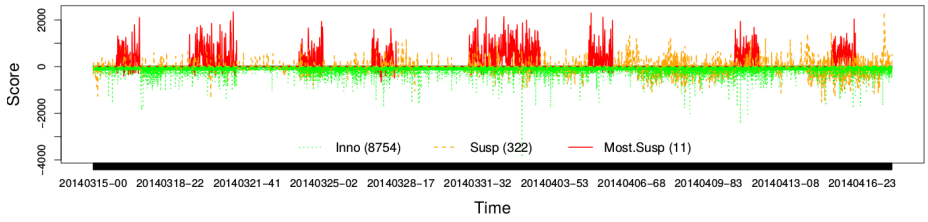

Figure 1: Monitoring from 15.03.2014 to 16.04.2014 (during the heartbleed public announcement).

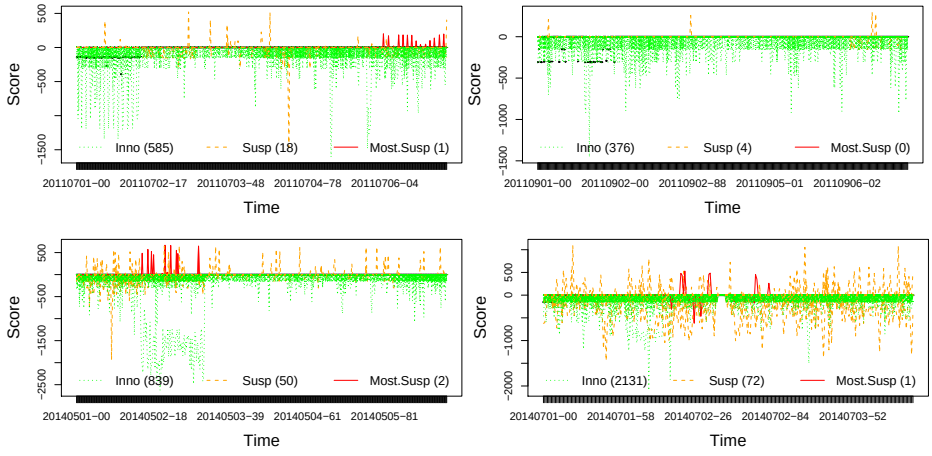

Figure 2: Monitoring from July 2011 to July 2014 (selected graphs).

Visual comparison between figures 1 and 2 gives an idea about the node behaviour over the period. Many nodes "stand out" from the normal behaviour during the heartbleed public announcement period in comparison to other periods. So, we would like to "early warn" about those nodes to carry out further investigations to classify their behaviour.

3 Research challenges

Ability to early warn depends upon three factors: the progression rate of attack lifecycle (e.g. a malware propagation gives more early warning time than a typical denial of service (DOS) attack), amount of evidence left at each stage, and the ability to acquire such evidence by sensors. This section highlights few challenges associated with these factors.

3.1 Generic set of indicators

In other domains such as natural disasters (e.g. tsunami), kinetic warfare and medical diagnosis (e.g. diabetes) early warnings are well established, and arguably simple when compared to early warnings on the cyberspace. For example, in kinetic warfare, intelligence officers study different sources of intelligence (e.g. listen to communications, satellite imagery) to looking for known preliminary indicators of military mobilisation. In medical diagnosis, preliminary indicators such as feeling thirsty, tired, losing weight and blurred vision early warn an individual about diabetes. But on the cyberspace, it is not clear what these indicators are or how they can be observed [6]. This presents a huge problem when trying to develop CEWS. As many scholars argue [7,6], CEWS cannot be developed from a purely technical perspective. They must consider more than just technical indicators and require significant input from other disciplines such as international relations and sociology as the focus of CEWS should be to warn of an impending attack rather than detecting when it in progress. However the biggest challenge, a generic set of indicators (signs) of preparation for an attack on the cyberspace is not well established (understood) yet.

3.2 Intelligence gathering

The cyberspace has a huge diversity. For example, it consists of different topological structures (e.g. PAN, LAN, MAN, WAN), different kind of networks (e.g. open Internet, darknet, honeynet, demilitarised zone) and different types of users (e.g. universities, health care system, the traffic system, power supply, trade, military networks). These entities produce events in different types and rates and have different analysis objectives and privacy requirements. In order to provide a representative image of the cyberspace at any given time, CEWS have to collect and process data from a range of these different entities. Employing a large monolithic sensor network for intelligence gathering on the cyberspace would not be possible due to these variations.

3.3 Uncertainty reasoning

The cyberspace is an uncertain place. Hence cyber defenders have to deal with a great deal of uncertainty [8,9] which is compounded by the nature of computing. Any future CEWS that seeks to model and reasoning on the cyberspace has to accept this ground truth and must deal with incompleteness (compensate for lack of knowledge), inconsistencies (resolve ambiguities and contradictions) and change (update the knowledge base over time). For example, entering misspelled password can be a simple mistake by an innocent user or a password guessing attempt by an attacker. Cyber defenders do not know who the attackers nor their location. Some suspicious events, e.g. a major router failure could generate many ICMP unreachable messages while some computer worms (e.g. CodeRed and Nimda) generate the same in active probing process, can appear as part of an attack as well as can originate from normal network activities. Other contextual information should be utilised to narrow down the meaning of such data [8].

3.4 Scalability

In principal it is possible to log every activity on every device on the cyberspace, but in practice security analysts cannot process these logs due to their vagueness as attack indicators as well as the sheer volume of data. The biggest challenge is how to start from imprecise and limited knowledge about attack possibilities, and quickly sift through huge volume of data to spot a small set of data that altogether makes the picture of attacks clear. As volume and rate of traffic are rising, inspection of each and every individual event is not feasible. A data reduction is needed [8].

3.5 Information fusion

As mentioned earlier, CEWS cannot be developed from a purely technical perspective. Given the huge number of possible data sources and overwhelming amount of data they generate, a data reduction method is essential to enable continuous security monitoring [10]. Future CEWS require fusing as many data sources as possible. Though it is not an exhaustive list some possible data sources for this task would be network data traffic, log files, social media, mobile location traces, mobile call traffic, web browsing traces, content popularity, user preferences, spatial/geographic distribution of network elements, network topology (router and AS level), network paths, protocol traces, social network structure and other security intelligence either system or social level.

3.6 Evaluation

Getting validity for a novel method is only possible through a proper evaluation. But in this research area, evaluation of novel algorithms against real time network data is a challenge. Real network traffic datasets with ground truth data on attack activity are difficult to obtain. Any such effort faces uncertainty of

success in investigating relevant patterns of activity. One solution to this problem would be to develop monitoring algorithms based on unary classification as it is relatively easier to find clean datasets than malicious ones, or providing mathematical proof for novel methods.

4 Related work

This section provides an overview of the existing practices for CEWS and provides a brief evaluation of some significant ideas to give future directions. Design and implementation of effective CEWS has a significant amount of overlaps with other research areas [6] such as situational awareness, intrusion detection and network monitoring. Hence we categorise them in related themes.

4.1 Threat scenario

Threat scenario provides an important aspect to the early warning discussion. For example, early warning on malware propagation can be easier than warning on DOS attack. The former needs a period of time to propagate and hence provides long early warning time (typically minutes to days). However early warning on DOS can be problematic as it might last within few seconds. Attempts to early warn on a particular threat type is common (e.g. [11,12,13,14,15]) in the literature. A malware warning centre is proposed in [11] which uses a Kalman filter to detect a worm's propagation at its early stage in real-time. An architecture of an automatic CEWS is discussed in [12]. Authors aim to provide predictions and advice regarding security threats without incorporation of cognitive abilities of humans. [13] aims for distributed, large-scale monitoring of malware on the Internet. A worm propagation stochastic model is built [14] to model the random effects during worm spreading by means of a stochastic differential equation. Authors propose a logical framework for a distributed early warning system against unknown and fast-spreading worms. An open-source early warning system to estimate the threat level and the malicious activities across the Internet is provided [15]. Limiting to a certain threat type is a major drawback of above proposals. They cannot simply extend for newly emerging threats.

4.2 Situational awareness

Situational awareness is an essential component of an CEWS, and hence related to this work. Cyber situational awareness includes awareness of suspicious network related activities that can take place at all levels in the TCP/IP stack [16]. Such activity can range from low-level network sniffing to suspicious linguistic content on social media. Various network measurements and techniques (e.g. packet inter arrival times [17], deep packet inspection [18], game theory [19]) have been employed in proposing these solutions. The idea for a common operational picture (big picture) is presented [20,21]. A systematic review of cyber

situational awareness can be found in [16]. However instead of addressing the full complexity, above solutions concentrated on a particular issue of the problem and some solutions (e.g. deep packet inspection) are neither feasible in practice nor suitable for real time analysis.

4.3 Information exchange

DShield internet storm centre is a cooperative network security community. It collects firewall and IDS logs world wide and incorporates human interpretation and action in order to generate predictions and advice [22]. eCSIRT.net [23] comprises of a sensor network which collects and correlates alerts for human inspection. The Internet motion sensor, a globally scoped Internet monitoring system aims to measure, characterise, and track threats [24]. It statistically analyses dark net traffic that needs to be interpreted by humans. DeepSight intelligence collects, analyses and delivers cyber-threat information through a editable portal and datafeeds, enabling proactive defensive actions and improved incident response [25]. Human analysis and data mining is incorporated in order to provide statistics. An infrastructure and organisational framework for a situation awareness and early warning system for the Internet is presented in [26]. This work aims for sharing, correlating and cooperatively analysing sensor data collected from number of organisations located in different geographical locations. eDare (Early Detection, Alert and Response system) [27] and the Agent-based CEWS [28] also focus on early warning in computer networks. However information exchange can be seen as a major barrier for CEWS' advances. In the context of security, data and information sharing is difficult between different organisations and nations due to various reasons [29,30]. An extensive survey of collaborative intrusion detection proposals can be found in [22].

4.4 Sensor networks

Sensing in-progress attacks requires strategically placed sensors throughout the cyberspace, and analysing acquired data to distinguish between attack traffic (events) and innocent traffic. Sensor network would be a vital part of CEWS. Current sensor networks for CEWS have a simple monolithic structure [31], where data is acquired at the network edges and then transmitted over a dumb infrastructure to a central location for analysis. This can cause various issues to the analysis due to many reasons such as nonidentical measurements, nonidentical local detectors and noisy channels [32]. High computational cost is another significant issue. Hence computationally fast and accurate methodology to evaluate the error, detection, and false alarm probabilities in such networks is essential. Optimal sensor placement strategies for CEWS is discussed in [33]. Authors study correlation between attack patterns of different locations (national and international) and explore how sensors should be located accordingly. The design and analysis of sensor networks for detection applications has received considerable attention during past decades [34].

4.5 Information fusion

Technical data itself is not sufficient to produce early warnings on computer networks. Fusion of different network measurements from different sources is essential. That measurement could be range from low-level network sniffing to suspicious linguistic contents on social media. A number of techniques have been employed last decades for information fusion on computer networks. Fusion of cyber-related information from a variety of resources including commercial news, blogs, wikis, and social media sources is proposed in [35]. Bayesian fusion for slow activity monitoring [36,8], high speed information fusion for real-time situational awareness [37], JDL data fusion model to computer networks [38], detecting network data patterns [39], combining data from sensors using ontology methods [40] and fuse security audit data with data from a psychological model [41] are few of them to mention. Using web-based text as a source for identifying emerging and ongoing attacks can be found in [42].

4.6 Tools and techniques

Most existing tools and techniques have been dedicated for security data analytic. An open, adaptable, and extensible visual analytic framework is provided in [43]. All data is treated as streaming and visualises them using machine learning techniques [44], live network situational awareness system that relies upon streaming algorithms included [45], fast calculations of important statistical properties of high speed and high volume data [45], sophisticated visualization of attack paths and automatic recommendations for mitigation [46] are some interesting works in the literature. In this context, there is a need to investigate "changes to the changing patterns" instead of changing points in a traffic profile sequence. This is essential as there are general systematic patterns (e.g. trend and seasonality) in the time series of user behaviours, and such variations should be considered as pretty normal in the analysis. Autocorrelations and differencing could help to deal with general dependencies in the data to make hidden patterns apparent and relevant.

5 Discussion

Instead of addressing the full complexity, existing works are concentrated on particular cyber issues such as sensor placement, type of sensors, data fusion and packet sampling. In order to deal with toady's advanced threats, an integrated large scale security analytic is needed. These advanced threats are the work of hacktivists, nation states, criminal enterprises and other groups with deep funding and specialised security expertise. They conduct reconnaissance not only on an organisation's security systems but also personnel and processes, and develop techniques to exploit them through social engineering, escalation of privileges and other forms of probing attacks. They move patiently through an organisation's network - taking days, weeks or months to accomplish their objectives - in

order to avoid detection [47]. In principal, early warning on such attacks is feasible as they provide long early warning time, but in practice research challenges discussed in the paper has to overcome to design and implement such a system.

As discussed in section 3.1 establishing a generic set of indicators is the biggest challenge. In empirical analysis in section 2, preliminary indicators defined using the existing knowledge of heartbleed attack. But how to derive a set of preliminary indicators for zero day attacks? One possible approach to address this would be building a complete (as much as possible) corpus of recently discovered attacks such as Stuxnet, Duqu 2.0 and Havex, and analysing that corpus in order to derive a generic set of indicators. This analysis should focus on each stage of attack lifecycle (e.g. reconnaissance, inspection paths, lateral movements, data ex-filtration and C2 activities) of each attack in the corpus.

Monolithic sensor network for intelligence gathering would not be suitable beyond research test beds. Deploying sensor networks with huge variations in administrative distribution and cooperation are required for advances in future CEWS. Investigations on how ordinary sensors can be employed to handle these type of complexities has not been covered much in the literature. Investigations to improve some exiting works (e.g. [32]) to fit this purpose would be interesting.

As discussed in section 3.3, not modelling the uncertainty in event classification is a major issue in many existing IDSs. As a result they produce huge number of false alarms, in which existing security monitoring tools bring significant amount of uncertainty to the true interpretation of security alerts. The uncertainty challenge exists in all stages of generic attack process [48]. There are three basics methods that can be employed to handle these kind of uncertainties: symbolic methods, statistical techniques and fuzzy methods. Efficient methods needed to leverage advances in these methods and other system level techniques for early estimation of malicious activities.

Applying a data reduction technique would be possible method to address scalability issues. Employing statistical sampling [49,8] and/or suitable approximation techniques (e.g. approximate Bayesian computation, saddle point approximation) would be possible methods to reduce the computational cost involved in the analysis. Node profiling through information fusion may address some issues such as storage [8]. Low-rank approximation [50] in minimisation problem in mathematical modelling and data compression would be interesting to investigate as a data reduction method on the cyberspace. Such a work can be found in [51].

As mentioned in section 3.5, analysing a centralised log collection or traffic capture is not longer enough for modern day security. While probabilistic fusion may be useful, a systematic investigation still needs to evaluate approaches for the ability to handle vagueness (fuzzy set), ambiguity (dempster-shafer) and incompleteness (possibilistic) of events, ultimately with an aim to develop hybrid data fusion techniques useful for early estimation. Events in the physical world offer additional sensors providing insight regarding the on going situation. Recent developments using Bayesian-based statistical profiling of potential targets of

cyber attacks provides for a promise to address this as it accommodates for analyst's prior belief [52,53].

References

1. Biskup, J., Hämmerli, B., Meier, M., Schmerl, S., Tölle, J., Vogel, M.: 2. 08102 working group-early warning systems. InProceedings biskup_et_al: DSP (2008) 1493
2. Lee, C., Yi, L., Tan, L.H., Goh, W., Lee, B.S., Yeo, C.K.: A wavelet entropy-based change point detection on network traffic: A case study of heartbleed vulnerability. In: Cloud Computing Technology and Science (CloudCom), 2014 IEEE 6th International Conference on. (Dec 2014) 995–1000
3. Durumeric, Z., Kasten, J., Adrian, D., Halderman, J.A., Bailey, M., Li, F., Weaver, N., Amann, J., Beekman, J., Payer, M., et al.: The matter of heartbleed. In: Proceedings of the 2014 Conference on Internet Measurement Conference, ACM (2014) 475–488
4. Cho, K., Mitsuya, K., Kato, A.: Traffic data repository at the wide project. In: Proceedings of the Annual Conference on USENIX Annual Technical Conference. ATEC '00, Berkeley, CA, USA, USENIX Association (2000) 51–51
5. Kalutarage, H., Shaikh, S., Lee, C., Sung, F.: Towards an early warning system for network attacks using bayesian inference. In: Cyber Security and Cloud Computing (CSCloud), 2015 IEEE 2nd International Conference on. (Nov 2015) 399–404
6. Robinson, M., Jones, K., Janicke, H.: Cyber warfare: Issues and challenges. Computers & Security **49** (2015) 70–94
7. Sharma, A., Gandhi, R., Mahoney, W., Sousan, W., Zhu, Q., et al.: Building a social dimensional threat model from current and historic events of cyber attacks. In: Social computing (SocialCom), 2010 IEEE second international conference on, IEEE (2010) 981–986
8. Kalutarage, H.K., Shaikh, S.A., Wickramasinghe, I.P., Zhou, Q., James, A.E.: Detecting stealthy attacks: Efficient monitoring of suspicious activities on computer networks. Computers & Electrical Engineering **47** (2015) 327–344
9. Kalutarage, H.: Effective monitoring of slow suspicious activites on computer networks. PhD thesis, Coventry University (2013)
10. Dempsey, K.: Information security continuous monitoring (ISCM) for federal information systems and organizations. US Department of Commerce, National Institute of Standards and Technology (2011)
11. Zou, C.C., Gao, L., Gong, W., Towsley, D.: Monitoring and early warning for internet worms. In: Proceedings of the 10th ACM conference on Computer and communications security, ACM (2003) 190–199
12. Apel, M., Biskup, J., Flegel, U., Meier, M.: Towards early warning systems–challenges, technologies and architecture. In: Critical Information Infrastructures Security. Springer (2010) 151–164
13. Engelberth, M., Freiling, F.C., Göbel, J., Gorecki, C., Holz, T., Hund, R., Trinius, P., Willems, C.: The inmas approach. (2010)
14. Magkos, E., Avlonitis, M., Kotzanikolaou, P., Stefanidakis, M.: Toward early warning against internet worms based on critical-sized networks. Security and Communication Networks **6**(1) (2013) 78–88
15. Kollias, S., Vlachos, V., Papanikolaou, A., Chatzimisios, P., Ilioudis, C., Metaxiotis, K.: Measuring the internet's threat level: A global-local approach. In: Computers and Communication (ISCC), 2014 IEEE Symposium on, IEEE (2014) 1–6

16. Franke, U., Brynielsson, J.: Cyber situational awareness–a systematic review of the literature. Computers & Security **46** (2014) 18–31
17. Harmer, P., Thomas, R., Christel, B., Martin, R., Watson, C.: Wireless security situation awareness with attack identification decision support. In: Computational Intelligence in Cyber Security (CICS), 2011 IEEE Symposium on. (April 2011) 144–151
18. King, D., Orlando, G., Kohler, J.: A case for trusted sensors: encryptors with deep packet inspection capabilities. In: Military Communication Conference, MILCOM 2012, IEEE (2012) 1–6
19. He, H., Xiaojing, W., Xin, Y.: A decision-support model for information systems based on situational awareness. In: Multimedia Information Networking and Security, 2009. MINES '09. International Conference on. Volume 2. (Nov 2009) 405–408
20. Cheng, Y., Sagduyu, Y., Deng, J., Li, J., Liu, P.: Integrated situational awareness for cyber attack detection, analysis, and mitigation. In: SPIE Defense, Security, and Sensing, International Society for Optics and Photonics (2012) 83850N–83850N
21. Preden, J., Motus, L., Meriste, M., Riid, A.: Situation awareness for networked systems. In: Cognitive Methods in Situation Awareness and Decision Support (CogSIMA), 2011 IEEE First International Multi-Disciplinary Conference on. (Feb 2011) 123–130
22. Zhou, C.V., Leckie, C., Karunasekera, S.: A survey of coordinated attacks and collaborative intrusion detection. Computers & Security **29**(1) (2010) 124–140
23. CSIRT_Network: The european computer security incident response team network. http://www.ecsirt.net/ (June 2015)
24. Bailey, M., Cooke, E., Jahanian, F., Nazario, J., Watson, D., et al.: The internet motion sensor-a distributed blackhole monitoring system. In: NDSS. (2005)
25. Symantec: Cyber security: Deepsight intelligence. http://www.symantec.com/deepsight-products/ (June 2015)
26. Grobauer, B., Mehlau, J.I., Sander, J.: Carmentis: A co-operative approach towards situation awareness and early warning for the internet. In: IMF. (2006) 55–66
27. Elovici, Y., Shabtai, A., Moskovitch, R., Tahan, G., Glezer, C.: Applying machine learning techniques for detection of malicious code in network traffic. In: KI 2007: Advances in Artificial Intelligence. Springer (2007) 44–50
28. Bsufka, K., Kroll-Peters, O., Albayrak, S.: Intelligent network-based early warning systems. In: Critical Information Infrastructures Security. Springer (2006) 103–111
29. Brunner, M., Hofinger, H., Roblee, C., Schoo, P., Todt, S.: Anonymity and privacy in distributed early warning systems. In: Critical Information Infrastructures Security. Springer (2011) 81–92
30. Koch, R., Golling, M., Rodosek, G.D.: Evaluation of state of the art ids message exchange protocols. In: International Conference on Communication and Network Security (ICCNS). (2013)
31. Theilmann, A.: Beyond centralism: The herold approach to sensor networks and early warning systems. In: Proceedings of First European Workshop of Internet Early Warning and Network Intelligence (EWNI 2010). (2010)
32. Aldosari, S., Moura, J.M., et al.: Detection in sensor networks: The saddlepoint approximation. Signal Processing, IEEE Transactions on **55**(1) (2007) 327–340
33. Göbel, J., Trinius, P.: Towards optimal sensor placement strategies for early warning systems. In: Sicherheit. (2010) 191–204
34. Varshney, P.K.: Distributed detection and data fusion. Springer Science & Business Media (1997)

35. Morris, T., Mayron, L., Smith, W., Knepper, M., Ita, R., Fox, K.: A perceptually-relevant model-based cyber threat prediction method for enterprise mission assurance. In: Cognitive Methods in Situation Awareness and Decision Support (CogSIMA), 2011 IEEE First International Multi-Disciplinary Conference on. (Feb 2011) 60–65
36. Chivers, H., Clark, J.A., Nobles, P., Shaikh, S.A., Chen, H.: Knowing who to watch: Identifying attackers whose actions are hidden within false alarms and background noise. Information Systems Frontiers **15**(1) (2013) 17–34
37. Sudit, M., Stotz, A., Holender, M.: Situational awareness of a coordinated cyber attack. In: Defense and Security, International Society for Optics and Photonics (2005) 114–129
38. Schreiber-Ehle, S., Koch, W.: The jdl model of data fusion applied x2014; a review paper. In: Sensor Data Fusion: Trends, Solutions, Applications (SDF), 2012 Workshop on. (Sept 2012) 116–119
39. Paffenroth, R., Du Toit, P., Nong, R., Scharf, L., Jayasumana, A.P., Bandara, V.: Space-time signal processing for distributed pattern detection in sensor networks. Selected Topics in Signal Processing, IEEE Journal of **7**(1) (2013) 38–49
40. Mathews, M.L., Halvorsen, P., Joshi, A., Finin, T.: A collaborative approach to situational awareness for cybersecurity. In: Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), 2012 8th International Conference on, IEEE (2012) 216–222
41. Greitzer, F.L., Frincke, D.A.: Combining traditional cyber security audit data with psychosocial data: towards predictive modeling for insider threat mitigation. In: Insider Threats in Cyber Security. Springer (2010) 85–113
42. Grothoff, K., Brunner, M., Hofinger, H., Roblee, C., Eckert, C.: " problems in web-based open source information processing for it early warning. (2011)
43. Jonker, D., Langevin, S., Schretlen, P., Canfield, C.: Agile visual analytics for banking cyber x201c;big data x201d;. In: Visual Analytics Science and Technology (VAST), 2012 IEEE Conference on. (Oct 2012) 299–300
44. Harrison, L., Laska, J., Spahn, R., Iannacone, M., Downing, E., Ferragut, E.M., Goodall, J.R.: situ: Situational understanding and discovery for cyber attacks. In: Visual Analytics Science and Technology (VAST), 2012 IEEE Conference on. (Oct 2012) 307–308
45. Streilein, W.W., Truelove, J., Meiners, C.R., Eakman, G.: Cyber situational awareness through operational streaming analysis. In: MILITARY COMMUNICATIONS CONFERENCE, 2011-MILCOM 2011, IEEE (2011) 1152–1157
46. Jajodia, S., Noel, S., Kalapa, P., Albanese, M., Williams, J.: Cauldron mission-centric cyber situational awareness with defense in depth. In: Military Communications Conference, 2011-MILCOM 2011, IEEE (2011) 1339–1344
47. Weber, D.: Transforming traditional security strategies into an early warning system for advanced threats. http://www.emc.com/collateral/software/solution-overview/h11031-transforming-traditional-security-strategies-so.pdf (September 2012)
48. Li, J., Ou, X., Rajagopalan, R.: Uncertainty and risk management in cyber situational awareness. In: Cyber Situational Awareness. Springer (2010) 51–68
49. Kalutarage, H.K., Shaikh, S.A., Zhou, Q., James, A.E.: Monitoring for slow suspicious activities using a target centric approach. In: Information Systems Security. Springer (2013) 163–168
50. Markovsky, I.: Low rank approximation: algorithms, implementation, applications. Springer Science & Business Media (2011)

51. Viswanath, B., Bashir, M.A., Crovella, M., Guha, S., Gummadi, K.P., Krishnamurthy, B., Mislove, A.: Towards detecting anomalous user behavior in online social networks. In: Proceedings of the 23rd USENIX Security Symposium (USENIX Security). (2014)
52. 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: Cyber Conflict (CYCON), 2012 4th International Conference on, IEEE (2012) 1–19
53. Shaikh, S.A., Kalutarage, H.K.: Effective network security monitoring: from attribution to target-centric monitoring. Telecommunication Systems (2015) 1–12