

EARLY DETECTION AND CONTAINMENT OF NETWORK WORM

^{1*} Muhammad Aminu Ahmad and ²Abdullahi Ismaila Jihad.
^{1,2}Department of Computer Science, Kaduna State University, Kaduna, Nigeria.

* Authors Email Addresses: ¹muhdaminu@kasu.edu.ng ²a.ismaila@kasu.edu.ng

ABSTRACT

This paper presents a network security framework for containing the propagation of network worms. The framework employs a detection mechanism at the network layer to identify the presence of a network worm and a data-link containment solution to block the infected host. A prototype of the mechanism has been used to demonstrate the effectiveness of the developed framework. An empirical analysis of network worm propagation has been conducted to test the framework. The results show that the developed framework is effective in containing network worms with almost no false positives.

Keywords: Containment, worm detection, malware, cyber defence.

1. INTRODUCTION

The Internet has provided a medium for communication and sharing of information amongst people, businesses, governments and organisations. Therefore, the Internet must be kept continuous and secured from any form of malicious activities such as unauthorised access to computer system and malware attacks. Malware can be classified under a number of headings, including viruses, worms, Trojans, spyware, adware, rootkits, drive-by downloads and other malicious and unwanted software (Jarno and Pirkka, 2013). Self-propagating malware, termed a worm, is a malicious software program that propagates across a network by infecting hosts and in some cases launching malicious activities. A scanning network worm propagates by probing pseudo-random addresses looking for vulnerable hosts, which makes the malware highly virulent in nature. Fast scanning network worms are a particularly dangerous sub-class of such software.

The Internet has experienced a number of notable worm outbreaks that caused disruption of services (Ahmad and Woodhead, 2015), damage to targeted systems (Falliere et al., 2011), cyber espionage (Boldizsár et al., 2012) and financial losses ranging from millions to billions of US Dollars (Craig, 2005). Common countermeasures to worm infection are signature-based antivirus software, network intrusion detection systems and host intrusion detection systems. However, the ability of such systems to counter the effects of fast scanning network worms is limited because their high propagation and infection rates pose a significant security threat, with consequent damage to networks and the Internet. Thus it is important to effectively identify and counter the propagation of worms, particularly fast scanning network worms, using a mechanism that must first work without the need to rely on signatures (sequence of byte-stream), because methods which rely on content signatures are unlikely to detect zero-day network worms. Additionally, due to the high propagation rates of fast scanning network worms, the traditional approach of waiting for patches to be released by vendors to fix vulnerabilities is not viable. This article presents a cross-layer detection and containment

technique as an improvement over the NEDAC mechanism reported by Ahmad and Woodhead (2015). The NEDAC mechanism uses datagram-header information at the network layer to detect the presence of fast scanning worms and a containment solution at data-link layer to block outgoing traffic from a host that has been identified as infected. The rest of the paper is organised as follows. Section 2 summarises related work on worm detection and containment systems. Section 3 presents a description of the NEDAC mechanism. Section 4 presents the evaluation method used to test the NEDAC mechanism. Section 5 presents and discusses the experimental results. Section 6 concludes this article and points out possible future work.

2. Related Works

A range of anomaly-based network intrusion detection systems have been developed to identify the presence of worms using datagram header information and payload information (Faisal et al., 2009).

Guofei et al. (2004) developed an algorithm, termed DSC, that correlates incoming and outgoing traffic, i.e., if a host received a datagram on port i , and then starts sending datagrams destined for port i , it becomes a suspect. Jung et al. (2004) proposed an algorithm, termed TRW, which identifies a remote host attempt to establish a new TCP connection to a local destination as normal if there is a corresponding TCP reply. On the other hand, failure to establish a successful TCP connection is considered suspicious. Nicholas et al. (2004) simplified the TRW scheme by considering all new connections to be a failure until a response is received. The algorithm drops a datagram if it does not match an existing and successfully-established connection after a predefined threshold count. These techniques only slow worm infections.

David et al. (2005) used DNS-based rate limiting to suppress scanning worms in an enterprise network by identifying the absence of DNS resolution before a new connection as anomalous. Shahzad and Woodhead (2014) proposed a scheme that uses the absence of DNS lookup action prior to an outgoing TCP SYN or UDP datagram to a new destination IP address to detect worm propagation, and a protocol termed Friends to spread reports of an identified worm event to potentially vulnerable and uninfected peer networks within the scheme. Li and Stafford (2014) proposed a worm detector, which they termed SWORD. SWORD comprises two main modules; a Burst Duration Detector (BDD) and a Quiescent Period Detector (QPD). The BDD module encompasses a burst detection algorithm to prevent fast scanning worms by creating a window for every different size of first-contact connections. The QPD module ensures that quiescent periods in network activity do not disappear because of constant worm scanning. These techniques consume resources in order to keep track of distinct connection and host information, especially in large networks (Pele et al., 2008) and they can only slow worm infections.

Additionally, Wang and Salvatore (2004) and Kim et al. (2012) proposed payload-based anomaly detection schemes. Wang

Salvatore (2004) proposed a detection scheme known as PAYL to detect and generate signatures for zero-day worms. PAYL uses a training phase to create a profile during normal operation, and produces a byte frequency distribution as a model for normal payloads. Based on this information, a centroid model is created and then during the detection phase, the Mahalanobis distance of each datagram payload from the centroid model is calculated. A datagram is considered to be anomalous based on its distance from the normal behaviour. Kim et al. (2012) proposed a detection scheme using a standalone device. The scheme employed the detection method reported by in Kim and Nnamdi (2008). During the training phase, the mean and standard deviation scores for all datagrams are computed. In the detection phase, a score is computed by counting the number of datagram bytes that fall outside the range defined for each byte. These above mentioned mechanisms have limitations such as computational complexity (Jyothsna, 2011), management overhead (Garcia-Teodoro et al., 2009), high rates of false positives (Pele et al., 2008) and incur significant delays in deployment and detection Kim et al. (2012).

3. The NEDAC Mechanism

The NEDAC mechanism comprises a network layer detection system and a containment system at the data-link layer that work together to provide a countermeasure solution, with a connection maintained between the two systems to enable continuous data transmission. The detection system detects anomalies from client hosts and server hosts in a network using different techniques. The detection system maintains a list of server IP addresses in order to differentiate client and server hosts in a network. Client hosts are defined as network hosts which typically consume Internet services (e.g. workstations, laptops, tablets, smart-phones, etc.) while server hosts are network hosts used to serve client requests (e.g. web servers and email servers).

The detection system keeps track of inbound and outbound TCP SYN and UDP datagrams for a window of time with value T to determine anomalies that exceed a threshold. The threshold is a maximum allowable count of anomalous datagrams a host can send before T has elapsed. The containment system receives the MAC address of an identified infected host from the detection system and then blocks all traffic originating from the host using MAC address access control. The working mechanism of NEDAC is presented in Listing 1.

Listing 1: The NEDAC Algorithm

```

11: Begin
12: /* Initialize tables */
13: initializeTable(inboundTable)
14: initilaizeTable(resolutionTable)
15: initilaizeTable(noResolutionTable)
16: /* initialize exempt table*/
17: initializeTable(exemptTable)
18: /* Set timer */
19: T = SetTimerSignal()
20: /* open interface */
21: openInterface(interface)
22: /* do in parallel */ process 1
23: while (there are datagrams to process) do
24:   getDatagram()
25:   uniqueMAC = getMACAddress(datagram)

```

```

16:   headerInfo = getHeaderInfo(srcIP,dstIP,sPort,dPort)
17:   if datagram is inbound then
18:     if packet is DNS reply then
19:       updateResolutionTable()
20:     else
21:       updateInboundTable(dstIP)
22:       updateInboundTable(dPort)
23:     end if
24:   else
25:     if headerInfo is not found in exemptTable[] then
26:       if source host not a server then
27:         if IP addresses not in resolutionTable then
28:           updateNoResolutionTable(headerInfo)
29:           if (dPortCounter > threshold) then
30:             containHost(MAC)
31:             if dPort in inboundTable then
32:               blockInbound(dPort)
33:             end if
34:           end if
35:         else
36:           getNextDatagram()
37:         end if
38:       else
39:         correlate(headerInfo)
40:         if dport in inboundTable then
41:           updateInboundTable(headerInfo)
42:         end if
43:         if (dPortCounter > threshold) then
44:           containHost(MAC)
45:           if dport in inboundTable then
46:             blockInbound(dPort)
47:           end if
48:         end if
49:       end if
50:     else
51:       getNextDatagram()
52:     end if
53:   end if
54: end while
55: /* do in parallel */ process 2
56: while true do
57:   if (T generates timeout signal) then
58:     for entries in resolutionTable do
59:       if (TTL >= 86400) then
60:         delete entry
61:       end if
62:     end for
63:     for entries in noResolutionTable and inboundTable do
64:       if (TTL >= 60) then
65:         delete entry
66:       end if
67:     end for
68:     halveThresholds()
69:   end if
70: end while
71: End

```

The NEDAC algorithm monitors TCP SYN and UDP datagrams from hosts in a network. For client hosts, the algorithm observes DNS resolution datagrams and records the IP address of the host that made the resolution and the resolved address in the resolution table. The destination IP address and port of an inbound datagram (excluding a DNS reply) is recorded in the inbound table for both client and server hosts. Outgoing datagram header information (source IP addresses and ports) is associated to entries in an exempt table. The exempt table comprises a list of IP addresses and ports that are exempt from the algorithm. If the header information results in a miss, the algorithm determines whether there is a recent DNS query by a client host for the destination IP address prior to sending the datagram by checking the resolution cache. If there is a miss, the algorithm records the destination port in the no-resolution cache, increments its counter and then determines excess using the threshold with value V . For server hosts, the algorithm checks the presence of the destination port of outbound TCP SYN and UDP datagrams in the inbound table. If there is a hit, a counter for such entry is incremented for TCP datagrams. An additional verification of the destination IP address is made to determine a reply UDP datagram, and if the destination IP address does not match the IP address recorded for such entry in the inbound table, its counter is incremented and then excess in threshold is also determined. Upon a host exceeding the set threshold by a server or client host, the algorithm invokes the containment system and then checks the presence of the suspect port in the inbound cache. If there is a hit, an additional countermeasure is applied at the network layer using an access control list (ACL) to block all inbound datagrams destined for the suspect port in the network segment. A time-to-live (TTL) is provided for entries in all the caches. The default TTL value for DNS (86400 seconds) is applied to the resolution cache and 60 seconds is applied to the no-resolution and inbound caches. Furthermore, the algorithm decrements the counters in the no-resolution and inbound tables by half after the expiration of a timing window of T , and then checks all caches to determine and remove entries with expired TTL values.

The improvements of the countermeasure mechanism on the previous technique (Ahmad and Woodhead, 2015) are (1) separate detection techniques for client and server hosts to improve effectiveness of the system (2) an additional countermeasure mechanism for inbound worm traffic to block remote to local worm infection and (3) time-to-live for records maintained in caches to reduce excessive resource consumption.

4. Evaluation Procedure

To evaluate the proposed mechanism, a software prototype was developed and tested using worm propagation experiments in a controlled environment. The NEDAC mechanism was tested along with two previously reported worm detection techniques namely DSC and DNS-based detection schemes. The schemes were also implemented in software based on the description provided by their authors. The DNS-based scheme was termed DNS-RL.

The testing environment used for the evaluation process is a virtualised testbed reported Ahmad et al. (2015). The testbed contains four virtualised enterprise networks comprising a number of virtual network cells. The testbed has a scale of 1200 virtual machines, supports the use of worm daemons and has utilities for

replaying network traces as background traffic. To generate background traffic during the worm propagation experiments, the evaluation used the DARPA 1999 evaluation dataset (Lippmann et al., 2000) The "inside" traces of weeks 1 and 3 of the dataset meet the requirements of the evaluation because they are attack free traces that contain payload information for the variety of protocols needed. Additionally, the traces include a wide range of collected traffic from 31 network hosts.

The evaluation process used two contemporary pseudoworms that were developed based on the Microsoft RDP (CVE-2012-0002) vulnerability (CVEa, 2014) of 2012 and the ShellShock vulnerability (CVEb, 2014) of 2014. Ahmad and Woodhead (2015) reported the likely susceptible population values and potential datagram sizes of the Microsoft RDP and ShellShock vulnerabilities as circa 16.5 M and 3800 bytes and 42.5 k and 2000 bytes respectively.

Additionally, the bandwidth available for an infected host and the worm datagram size determine how fast a worm can send datagrams. The average Internet connection speed was estimated to be within the range 10 Mbps to 1000 Mbps (NetIndex, 2014). Although it is impossible for a host to achieve the maximum speed of a network card, the vast majority of Internet connected hosts are capable of transmitting data at 60 Mbps to 120 Mbps (Marshini et al., 2011). Thus based on the assumption that the Internet connected hosts exhibit an average data transmission rate of 90 Mbps, the scan rate S , required for a single worm instance to transmit a datagram of size M (in bytes), over a C megabits Internet connection per second can be determined using equation 1.

$$s = \frac{C}{M * 8} \quad (1)$$

Therefore, the likely scan rates for the RDP and ShellShock pseudo-worms are $\left(\frac{9,000,000}{3800 * 8}\right) = 2960$ and $\left(\frac{9,000,000}{2000 * 8}\right) = 5625$ datagrams per second respectively.

The scan rates of the pseudo-worms were scaled down by a factor of 24 and 45 for the RDP and ShellShock pseudoworms respectively to avoid overloading server resources. The resulting scan rates employed in the experiments are 125 "infectious" datagrams per second for RDP and ShellShock. Furthermore, the results of the experiments were scaled up by a factor of 24 and 45 for the RDP and ShellShock pseudoworms respectively.

Ahmad and Woodhead (2015) reported the number of susceptible hosts per million Internet hosts for RDP and ShellShock pseudo-worms as 4454 and 12 respectively. Thus, due to the scale of the testbed used, which has a maximum number of 1200 hosts, four class B size 2^{16} networks were used for RDP and five class A size 2^{24} networks were used for ShellShock. Thus, the resulting values were $\left[2^{16} * 4 * \left(\frac{4454}{1,000,000}\right)\right] = 1168$ and $\left[2^{24} * 5 * \left(\frac{12}{1,000,000}\right)\right] = 1007$ susceptible hosts respectively, within the relevant network address space.

A. Experimentation Setup

The evaluation experiments were conducted using the software prototypes of NEDAC DSC and DNS-RL. During the evaluation, a prototype of a detection scheme was positioned on the gateways

of each network, and for NEDAC, the containment system was positioned on the switches as depicted in Fig. 1.

The RDP and ShellShock worm propagation experiments were conducted using random and then hit-list scanning behaviours for each detection scheme. The random scanning technique probes IPv4 addresses within the routable address space. The hit-list scanning technique infects a list of pre-compiled vulnerable hosts and then each infected host uses random scanning. For each pseudo-worm experiment, a number of hosts (1168 for RDP and 1007 for ShellShock) were configured with the correct daemon to make them vulnerable to worm attack datagrams while other hosts were configured to replay the DARPA traces as background traffic. The worm attack and traffic replay events were executed concurrently in each experiment. The experiments were conducted without any countermeasures in place, then repeated with the countermeasures and the DARPA dataset as background traffic using threshold values of 100 through 400 anomalous datagrams sent by a host in a timing window of 10 seconds. The worm infection event was initiated by sending a UDP datagram to one of the vulnerable hosts.

RDP Pseudo-worm

The RDP pseudo-worm experiment was conducted using 1160 client hosts and 8 server hosts. The pseudo-worm daemon was configured to listen on UDP port 3389 and then transmit UDP datagrams to port 3389 at a scan rate of 125 "infectious" datagrams per second, once "infected". Five RDP pseudo-worm experiments were conducted using one initially infected host. The RDP-based worm experiment was repeated with a hitlist [25] of 10 and 20 hosts.

ShellShock Pseudo-worm

The ShellShock pseudo-worm experiment was conducted using 996 client hosts and 10 server hosts. The pseudo-worm daemon was configured to listen on UDP port 8080 and then transmit UDP datagrams to port 8080 at a scan rate of 125 "infectious" datagrams per second, once "infected". Five ShellShock pseudo-worm experiments were conducted using one initially infected host. As with RDP, the ShellShock worm experiment was repeated with a hit-list (Staniford, 2002) of 10 and 20 hosts.

5. RESULTS

This section discusses the results of the experiments conducted

using the candidate pseudo-worms for random and hit-list scanning and the false positives observed during the experiments. The results of the worm propagation and detection performance are presented in Fig. 2 through Fig. 9.

A. Random Scanning Infection

The results of random infection behaviours for the RDP and ShellShock pseudo-worms using a threshold value of 100 are presented in Fig. 2 and Fig. 5. When no countermeasure solution was in place, the RDP pseudo-worm infected 95% (1110) of the hosts in eight seconds as shown in Fig. 2. Additionally, the ShellShock pseudo-worm infected 95% (956) of its susceptible hosts in 145 seconds as shown in Fig. 5. When the detection schemes were applied, the infections were delayed and suppressed by DSC and DNS-RL and blocked completely by NEDAC. With DSC and DNS-RL, the RDP pseudo-worm infection was delayed by 12 seconds and suppressed to 44% (510) and 50% (580) respectively. The worm infections were detected by the DSC and DNS-RL schemes and the countermeasure solution was applied, but the initially infected host continued sending infectious datagrams, which infected a large number of hosts. However, with NEDAC, the initially infected host, for each pseudo-worm experiment, was detected and then blocked from sending out datagrams at the data-link layer, which stopped the infection completely for each of the two worm outbreak scenarios.

Additionally, the NEDAC scheme blocked inbound traffic destined for the destination port used by the identified worm infection at the network layer, which also enable the mechanism to contain the worm infection quickly.

B. Hit-list Scanning Infection

Fig. 3 and Fig. 6 show the results of the worm experiments conducted with a hit-list of 10 hosts. When no countermeasure was in place, the RDP pseudo-worm infected 95% (1110) of the hosts in 6 seconds as shown in Fig. 3. The ShellShock pseudo-worm attained 95% (956) infection in 55 seconds as shown in Fig. 6. With the DSC and DNS-RL scheme, the RDP pseudoworm infection attained 95% in 19 and 15 seconds respectively

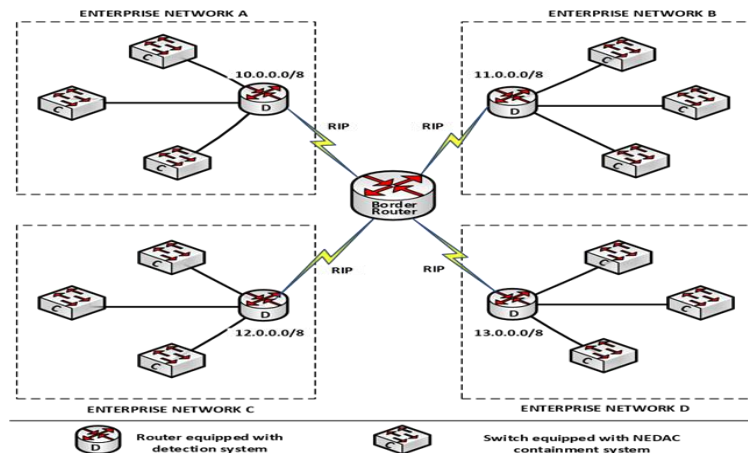


Fig. 1: Prototype setup




Fig. 2: RDP random scanning

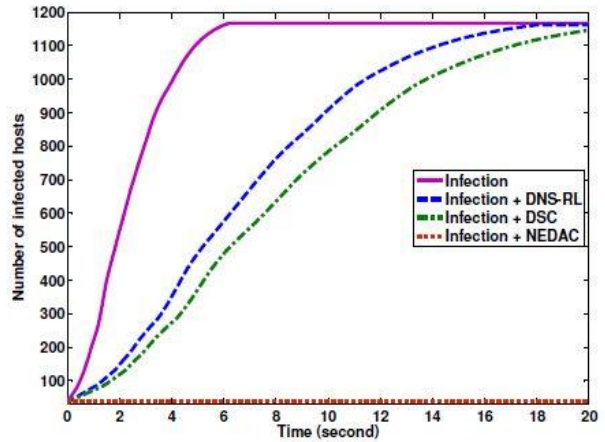


Fig. 3: RDP with a hit-list of 10 hosts

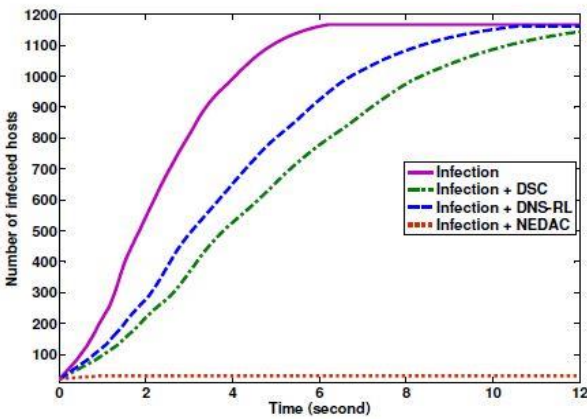


Fig. 4: RDP with a hit-list of 20 hosts

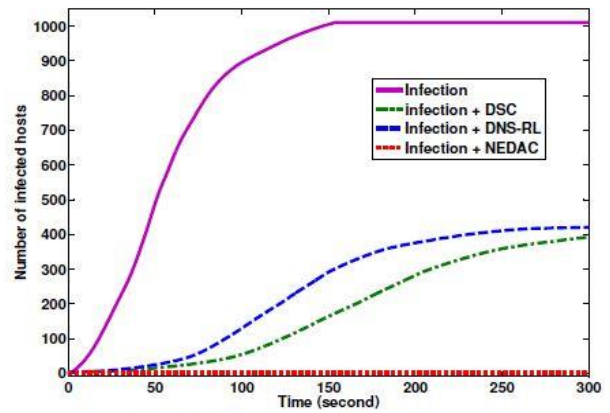


Fig. 5: ShellShock random scanning




Fig. 6: ShellShock with a hit-list of 10 hosts




Fig. 7: ShellShock with a hit-list of 20 hosts

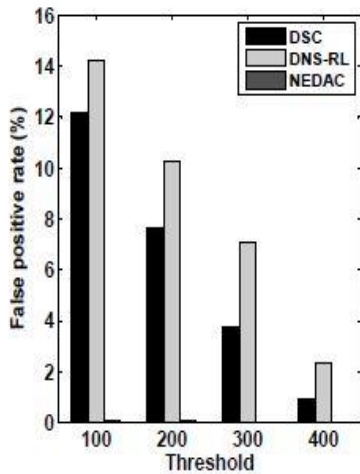


Fig. 8: False positive rate for RDP experiment




Fig. 9: False positive rate for ShellShock experiment

The ShellShock pseudo-worm attained 95% infection in 150 and 100 seconds with DSC and DNS-RL respectively. Furthermore, nine further infections were observed with NEDAC during the RDP pseudo-worm propagation and no further infections were observed during the propagation of the ShellShock pseudo-worm. Fig. 4 and Fig. 7 show the results of the worm experiments conducted with a hit-list of 20 hosts. When no countermeasure was in place, the RDP pseudo-worm infected 95% (1004) of the hosts in 5 seconds as shown in Fig. 4. The ShellShock pseudo-worm attained 95% infection in 40 seconds as shown in Fig. 7. Furthermore, with the DSC and DNS-RL scheme, the RDP pseudo-worm infection attained 95% in 11 and 9 seconds respectively. The ShellShock pseudo-worm attained 95% infection in 90 and 75 seconds with DSC and DNSRL respectively. For NEDAC, 56 further infections were observed during the RDP pseudo-worm propagation and no further infections were observed during the propagation of the ShellShock pseudo-worm.

C. Detection Performance

The false positive rates observed by the three detection schemes are presented in Fig. 8 and Fig. 9 for the RDP and ShellShock pseudo-worm experiments respectively. The schemes detected all real pseudo-worm datagrams in all the experiments conducted and therefore the true positive (TP) rates are 100%. However, the DSC and DNS-RL schemes incurred higher rates of false positives (FP) than NEDAC. NEDAC has very low FP rates using 100 and 200 as thresholds and zero FP rates using 300 and 400 as thresholds. Generally, the false positive rate diminishes with rising threshold values. NEDAC raised one false positive with threshold values of 100 and 200, which was caused by a multicast UDP datagram sent by a host to port 520, that is, a RIP advertisement. Additionally, the rate at which RIP sends updates to neighbouring routers is not similar to fast scanning worm behaviour because RIP routers exchange update every 30 seconds by default. Nevertheless, the RIP port can be added into the exempt list in NEDAC to avoid false positives. Across the whole experimental data set, NEDAC has a better performance in terms of false positives compared to the DNS-RL scheme.

6. Conclusion and Future Work

This paper has presented a countermeasure solution against fast scanning network worms. A software prototype of the worm countermeasure solution was used to evaluate the scheme using a set of experiments. The results of the experiments showed that the countermeasure solution is sensitive in detecting and containing an identified worm infection with almost no false positives. The results of a comparative analysis showed that the countermeasure solution has a better performance compared to two previously reported detection schemes.

As for future work, it is desirable to evaluate the mechanism using different background traffic. The aim of improving the detection system to use dynamic threshold policy and the speed of containment will also be investigated. Furthermore, the effect of timing window size and volume of background traffic will be investigated

REFERENCES

- Ahmad M. A. and Woodhead S. (2015) *Containment of Fast Scanning Computer Network Worms*. In Internet and Distributed Computing Systems, volume 9258 of Lecture Notes in Computer Science, pages 235–247. Springer International Publishing, 2015.
- Ahmad M. A., Woodhead S. and Diane Gan. *The V-Network Testbed for Malware Analysis*. In International Conference on Advanced Communication Control & Computing (ICACCCT), May 2016.
- Boldizsár B., Gábor P., Levente B. and Mark F. (2012) *The Cousins of Stuxnet: Duqu, Flame, And Gauss*. Future Internet, 4(4):971–1003, 2012.
- Craig F. (2005) *Computer worms: past, present, and future*. East Carolina University, 8, 2005.
- CVEa (2014) *Common Vulnerabilities and Exposures*. [Online]. Accessed on 19th October 2014. <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2012-0002>.
- CVEb (2014) *Common Vulnerabilities and Exposures*. [Online]. Accessed on 19th October 2014. <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2014-0002>.

6271.

- David W., Evangelos K. and Paul C. V. (2005) *DNS-Based Detection of Scanning Worms In An Enterprise Network*. In NDSS, 2005.
- Faisal M. C., Adeel A. and Zeshan I. (2009) *Comparative Evaluation of Header Vs. Payload Based Network Anomaly Detectors*. In Proceedings of the World congress on Engineering, 1: 1–5.
- Falliere N., Murchu L. O., and Eric C. (2011) W32. Stuxnet Dossier. White paper, Symantec Corp., Security Response, 5, 2011.
- Garcia-Teodoro P., Diaz-Verdejo J., Maciá-Fernández G. and Vázquez E. (2009) *Anomaly-Based Network Intrusion Detection: Techniques, Systems and Challenges*. computers & security, 28(1):18–28, 2009.
- Guofei G., Monirul S., Xinzhou Q., David D., Wenke L. and George R. (2004) *Worm Detection, Early Warning and Response Based On Local Victim Information*. In Computer Security Applications Conference, pages 136–145. IEEE, 2004.
- Jarno N. and Pirkka P. (2013) *Malware Detection by Application Monitoring*. US Patent 8,590,045, November 19 2013.
- Jung J., Paxson V., Berger A. W. and Balakrishnan H. (2004) *Fast Portscan Detection Using Sequential Hypothesis Testing*. In Security and Privacy, 2004. Proceedings. 2004 IEEE Symposium on, pages 211– 225. IEEE, 2004.
- Jyothsna V., Rama P. V. and Prasad K. M. (2011) *A Review of Anomaly Based Intrusion Detection Systems*. International Journal of Computer Applications, 28(7):26–35, 2011.
- Kim S. and Nnamdi N. (2008) *Noise-Resistant Payload Anomaly Detection for Network Intrusion Detection Systems*. In Performance, Computing and Communications Conference, 2008. IPCCC 2008. IEEE International, pages 517–523. IEEE, 2008.
- Kim S., Nnamdi N., William E., Blake J. and Paloma F. (2012) *On Network Intrusion Detection for Deployment in The Wild*. In Network Operations and Management Symposium (NOMS), 2012 IEEE, pages 253–260. IEEE, 2012.
- Li J. and Stafford S. (2014) *Detecting Smart, Self-Propagating Internet Worms*. In Communications and Network Security (CNS), 2014 IEEE Conference on, pages 193–201. IEEE, 2014.
- Lippmann R, Haines J. W., Fried D. J., Korba J. and Kumar Das. *The 1999 Darpa Off-Line Intrusion Detection Evaluation*. Computer networks, 34(4):579–595, 2000.
- NetIndex (2014) [Online]. Accessed 16th November 2014. Available: <http://www.netindex.com/>.
- Nicholas W., Stuart S. and Vern P. (2004) *Very Fast Containment of Scanning Worms*. In USENIX Security Symposium, 2: 16–85, 2004.
- Marshini C., David H., Andrew B., Ugochi O., Bethany S. and Rebecca G. (2011) *Why Is My Internet Slow? Making Network Speeds Visible*. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, pages 1889–1898. ACM, 2011.
- Pele L., Mehdi S. and Xiao S. (2008) *A Survey of Internet Worm Detection and Containment*. Communications Surveys & Tutorials, IEEE, 10(1):20– 35, 2008.
- Shahzad K. and Woodhead S. (2014) *Towards Automated Distributed Containment of Zero-Day Network Worms*. In Computing, Communication and Networking Technologies (ICCCNT), 2014 International Conference on, pages 1–7. IEEE, 2014.
- Wang K and Salvatore J. S. (2004) *Anomalous Payload-Based Network Intrusion Detection*. In Recent Advances in Intrusion Detection, pages 203–222. Springer, 2004.
- Staniford S., Paxson V. and Weaver N. (2002) *How to Own the Internet in Your Spare Time*. In USENIX Security Symposium, pages 149–167, 2002.