WHO'S NEXT? IDENTIFYING RISK FACTORS FOR SUBJECTS OF TARGETED ATTACKS

Martin Lee
Symantec.cloud, 1240 Lansdowne Court,
Gloucester, GL3 4AB, UK

Email martin_lee@symantec.com

ABSTRACT

Malware-containing emails can be sent to anyone. Single malware variants can be sent to tens of thousands of recipients without distinction. However, a small proportion of email malware is sent in low copy number to a small set of recipients that have apparently been specifically selected by the attacker. These targeted attacks are challenging to detect and if successful, may be particularly damaging for the recipient.

The vast majority of Internet users will never be sent a targeted attack. The few users to which such attacks are sent, presumably possess features that have brought them to the attention of attackers, and have caused them to be selected for attack. Applying epidemiological techniques to calculate the odds ratio for features of malware recipients, both targeted and non-targeted, allows the identification of putative factors that are associated with targeted attack recipients.

In this paper we show that it is possible to identify putative risk factors that are associated with individuals subjected to targeted attacks by considering the threat akin to a public health issue. These risk factors may be used to identify those at risk of being subject to future targeted attacks, so that these individuals can take additional steps to secure their systems and data.

INTRODUCTION

In recent years targeted malware attacks have frequently made press headlines. Low copy number, sophisticated malware sent by attackers to a small set of recipients is particularly difficult to detect and to protect against. The incidence of such attacks appears to be increasing. The *Symantec.cloud* mail-scanning service reports that an average of 77 attacks per day were detected in 2010, rising to 82 per day during 2011 [1]. However, these figures must be considered in the context of approximately 500,000 pieces of malware being blocked per day by the same service.

Information security officers are justified in being concerned about such attacks – successful attacks can be enormously expensive. One recent high-profile targeted attack has reportedly cost the breached organization $66 million in direct costs alone [2, 3]. Protecting against such threats may require dedicated resources and long-term investment [4]. Nevertheless, in an environment of static or decreasing budgets, information security departments may be reluctant to invest in protection against a rare threat that may never be encountered.

Currently, organizations that wish to understand their degree of exposure to such attacks are poorly served by the lack of techniques and data from which informed decisions can be made. Forensic tools may be able to report in retrospect that an organization has been subjected to targeted attacks, but cannot report the degree of risk that the organization faces. Identifying whether an organization is likely to be targeted, and which individuals within the organization are most likely to receive an attack, allows resources to be appropriately allocated. To achieve this level of understanding of relative risks, we must research the process that leads to individuals becoming the subject of a targeted attack, and collect empirical data.

To develop the necessary tools to assist with the calculation of the risk of targeted attacks, it is helpful to look elsewhere at other problem domains to see how similar issues have been addressed. In many ways, targeted attacks are akin to a public health issue. Individuals subject to targeted attacks tend to share many characteristics, such as working in certain industries, or having a certain level of seniority [1]. Considering such characteristics as possible risk factors for being subject to attack and calculating their significance can lead to the identification of individuals who are at high risk of attack and the identification of behaviour that is associated with being attacked.

Epidemiology

Epidemiology is the science concerned with the '*incidence, distribution and control of disease in a population*'[5]. Dr John Snow is recognized as the father of modern epidemiology. During the 1854 outbreak of cholera in London, he identified that living in proximity to a public water source in Broad Street was associated with the disease. Research of the exceptions to this observation supported the idea that drinking water from the pump caused the disease: the inhabitants of the nearby workhouse did not contract cholera, but they had their own separate water supply. A widow living some miles away did contract the disease, but, liking the taste of Broad Street water, she had her sons bring her a supply [6].

Almost a century later, Sir Richard Doll and Austin Bradford Hill discovered the link between smoking and lung cancer by questioning British physicians regarding their smoking habits and following their subsequent health and mortality [7, 8].

These epidemiological techniques developed to associate lifestyle factors with adverse health outcomes can be applied to information security. In place of 'adverse health outcome' we can substitute 'security incident', and use the same technique of correlating behavioural and lifestyle factors to these incidents. In this way we can identify the putative factors that predispose individuals to becoming affected by security incidents and intervene earlier to reduce exposure and minimize the consequences.

The Centre for the Protection of National Infrastructure describes the motivations behind emails containing targeted trojans as follows: '*The attackers' aim appears to be covert gathering and transmitting of commercially or economically valuable information*' [9]. We can hypothesize that individuals with access to 'commercially or economically valuable information' may be at risk of being subject to targeted attacks.

The question is how to identify which workers with access to what information are most at risk, so that these individuals can be better protected.

Odds ratio

One epidemiological technique that may be applied to information security is the calculation of odds ratio [10]. This technique allows researchers to work backwards from a population afflicted with an adverse outcome, and to compare them to a similar, but unafflicted control population, in order to identify factors that are associated with the adverse outcome. In this type of case-control analysis it is vital to ensure that the unafflicted control group is as similar as possible to the afflicted group to avoid bias [11].

The correlation of putative risk factors with the adverse outcome can be discovered by calculating the following table:

	Afflicted	Unafflicted
With risk factor	p_{11}	p_{10}
Without risk factor	p_{01}	p_{00}

Where p_{11} is the probability of afflicted individuals possessing the risk factor (i.e. the number of afflicted individuals possessing the risk factor divided by the total number of afflicted individuals).

Conversely, p_{01} is the probability of afflicted individuals not possessing the risk factor.

p_{10} is the probability of unafflicted individuals within the control group also possessing the risk factor, and p_{00} is the probability of unafflicted individuals in the control group not possessing the risk factor.

The odds ratio (OR) is calculated as:

$$\text{OR} = \frac{p_{11}\, p_{00}}{p_{10}\, p_{01}}$$

An odds ratio >1 implies a positive correlation for the risk factor, that the risk factor is more likely to be found in the afflicted group than the unafflicted group. An odds ratio <1 implies a negative correlation; the risk factor is less likely to be found in the afflicted group than the unafflicted. In this case the factor may be thought of as a protective factor.

The standard error for the natural logarithm of the odds ratio can be calculated as:

$$\text{SE}(\log_e \text{OR}) = \sqrt{\frac{1}{n_{11}} + \frac{1}{n_{10}} + \frac{1}{n_{01}} + \frac{1}{n_{00}}}$$

Where n_{11} is the number of afflicted individuals possessing the risk factor, n_{10} is the number of afflicted individuals without the risk factor, n_{01} is the number of control unafflicted individuals with the risk factor, n_{00} is the number of control unafflicted individuals without the risk factor.

The upper and lower 95% confidence values (W,X) for the natural logarithm of the odds ratio are calculated as:

$$W = \log_e \text{OR} - (1.96\ \text{SE}(\log_e \text{OR}))$$

$$X = \log_e \text{OR} + (1.96\ \text{SE}(\log_e \text{OR}))$$

The 95% confidence interval for the odds ratio is the exponential of W and X, e^w to e^x. That is to say to be 95% certain that the risk factor is positively correlated, both e^w and e^x should be greater than 1; for the risk factor to be likely to have a negative correlation, both e^w and e^x should be less than 1 [12].

METHODS

Experimental dataset

Symantec collects data regarding targeted attacks that consist of emails with malicious attachments. These emails are identified from the vast majority of non-targeted malware by evidence of there being prior research and selection of the recipient, with the malware being of high sophistication and low copy number. The process by which the *Symantec.cloud* mail scanning service collects such malware has already been described elsewhere [13, 14]. Other forms of targeted attack where the malicious payload is not attached to an email are probably associated with attack campaigns, but are not included in the dataset. The corpus almost certainly omits some attacks, and most likely also includes some non-targeted attacks, but nevertheless it represents a large number of sophisticated targeted attacks compiled according to a consistent set of criteria which render it a very useful dataset to study.

Experimental design

A reasonable hypothesis is that targeted attacks may be associated with the area of work of the recipients. If this is the case, we would expect to see a significant correlation between certain work domains and being subjected to targeted attacks compared with individuals receiving non-targeted attacks.

In order to calculate the correlation between work subject and targeted attack, we require a means of recording the subject of work of individuals and a population who receive targeted attacks for whom we can ascertain their subject of work.

The dataset of targeted attacks includes the email address of the recipient. Often, information regarding the nature of the work of the recipient is available from online sources, or business social networking websites. However, frequently the information is vague, or the individual has changed jobs to work in different sectors, which frustrates categorization.

Researchers in academic institutions present many advantages as a study group. Researchers are characterized by the fact that they publish their research, therefore their area of expertise and work is easy to ascertain. Additionally, many researchers have personal home pages where they list their recent publications along with the faculty and department in which they work.

Comprehensive ontologies for categorizing academic subjects are also available. The Joint Academic Coding System (JACS) developed by the Higher Education Statistics Agency is a

classification scheme for higher education courses offered in the UK [15]. This scheme covers all academic disciplines and offers two levels of granularity, a short code and a more extensive long code, to which academic research subjects can be mapped.

Many jurisdictions offer distinctive second-level domain extensions by which academic email addresses may be identified, in addition to the top-level '.edu' domain for accredited US educational institutions. These '.ac.' and '.edu.' domain types may be used not only to identify academic recipients of targeted attacks within the targeted attack corpus, but also to identify suitable control subjects in datasets of non-targeted malware attacks.

The research activities of the 182 academic recipients of targeted attacks between January 2010 and December 2011 were mapped to the JACS ontology. Each recipient was assigned to a single JACS code that was judged as best matching their research interests. JACS is designed primarily as a classification of undergraduate degree subjects, therefore some subjects, such as postgraduate medical subjects, are not assigned a distinct JACS code. In these cases, the code that was judged most representative either from biological sciences, or the code for clinical medicine was used.

Only the fact that the recipient had received an attack during the study period was counted. Many recipients received more than one attack, however investigating the association between frequency of attack and risk factors was not a goal of the study.

The recipients were selected as possessing an email address domain containing '.ac.' or '.edu.' or ending in '.edu'. One recipient was an alumni address which was excluded from the study, three recipient addresses were identified synonyms for other addresses and were considered as attacks against single individuals rather than separate attacks.

Additional classifications of 'staff' for non-academic employees of academic institutions, 'unknown' for email addresses where no information could be found regarding the recipient, and 'mailbox' for shared email addresses were used in addition to the JACS codes.

A control group of 188 academic recipients of non-targeted Bredolab email malware were randomly selected. The same email address pattern for targeted attack recipients was used to identify recipients. Classification of research interests was performed identically.

Odds ratios for each JACS short subject code and additional classifications were calculated according to Table 1.

RESULTS

Using the JACS ontology of short subject codes results in a total of 19 subject codes being found in both the targeted and control groups. Three individuals classified as working within the area

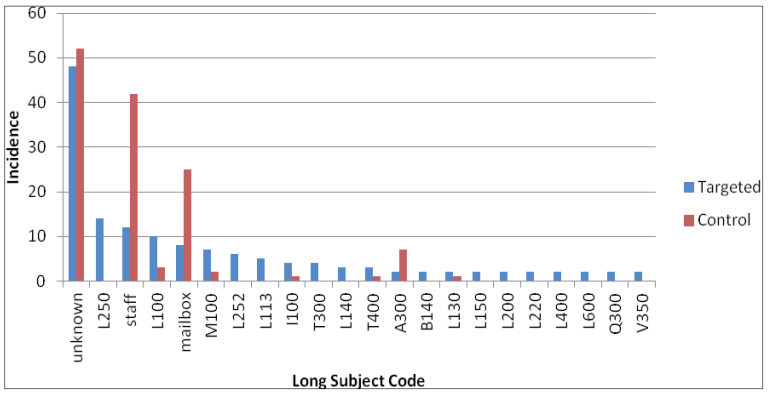

Graph 1: Incidences of most targeted long subject codes for the targeted and control groups.

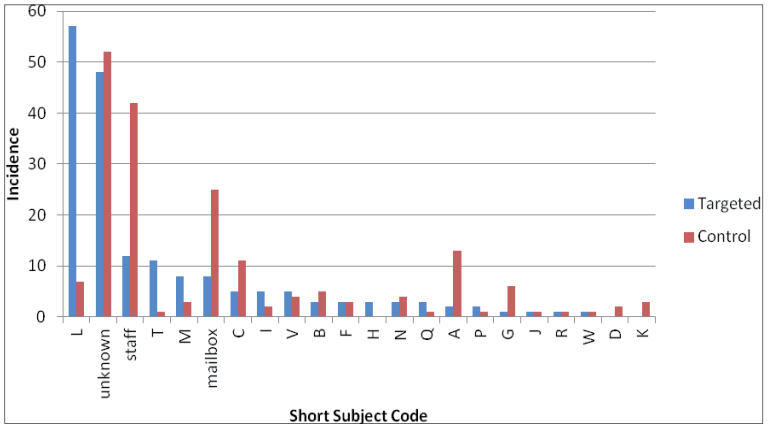

Graph 2: Incidences of all short subject codes for the targeted and control groups.

	Received a targeted attack email (n_0)	Received a non-targeted attack malware email (n_1)
Classified with code	p_{11}	p_{10}
Not classified with code	p_{01}	p_{00}

$n_0 = 182$, $n_1 = 188$

Table 1: Calculation of odds ratios.

of Engineering, subject code 'H', received targeted attacks, but no control subjects were counted within this subject. Therefore the odds ratio of this subject could not be calculated. Conversely, Veterinary Science, Agriculture and related subjects, and Architecture, Building and Planning received no targeted attacks but were found in the control group, giving odds ratios of 0. Education, subject code 'X', received no targeted attacks and was not found in the control group.

Short subject codes 'L', Social Studies, and 'T', Eastern, Asiatic, African, American and Australasian Languages,

Subject Code	Subject	Odds Ratio	95% Confidence Interval
A	Medicine & Dentistry	0.15	(0.03 – 0.67)
B	Subject Allied to Medicine	0.61	(0.14 – 2.60)
C	Biological Sciences	0.45	(0.15 – 1.34)
D	Veterinary Science, Agriculture and Related Subjects	0	-
F	Physical Sciences	1.03	(0.21 – 5.19)
G	Mathematical Sciences	0.17	(0.02 – 1.41)
I	Computer Sciences	2.63	(0.50 – 13.72)
J	Technologies	1.033	(0.06 – 16.64)
K	Architecture Building & Planning	0	-
L	**Social Studies**	**11.79**	**(5.21 – 26.70)**
M	Law	2.83	(0.74 – 10.86)
Mailbox		0.300	(0.13 – 0.68)
N	Business & Administrative Studies	0.77	(0.17 – 3.49)
P	Mass Communication & Documentation	2.08	(0.19 – 23.12)
Q	Linguistics, Classics and Related Subjects	3.13	(0.32 – 30.41)
R	European Languages, Literature and Related Subjects	1.03	(0.06 – 16.64)
Staff		0.25	(0.12 – 0.48)
T	**Eastern, Asiatic, African, American and Australasian Languages, Literature and Related Subjects**	**12.03**	**(1.54 – 94.16)**
Unknown		0.94	(0.59 – 1.48)
V	Historical and Philosophical Studies	1.30	(0.34 – 4.92)
W	Creative Arts and Design	1.03	(0.06 – 16.64)

Table 2: Odds ratio and 95% confidence interval for subject codes.

Literature and Related Subjects, are both positively correlated with targeted attacks at more than 95% confidence.

Within social studies, the most targeted long subject codes are 'L100' and 'L250', Economics and International Relations respectively. 'L100' on its own narrowly fails to reach the criteria for being positively correlated with targeted attacks with an odds ratio of 3.59, 95% confidence interval (0.97 – 13.25). 'T300', South Asian Studies, is the most attacked long subject code within group T. However, without this subject occurring within the control group, it is not possible to calculate an odds ratio.

Subject codes 'A', Medicine & Dentistry, 'D', Veterinary Science, Agriculture and Related Subjects, 'K', Architecture Building & Planning, shared mailbox addresses and non-academic staff are all negatively correlated with targeted attacks. Such classifications can be thought of as protective factors in that email addresses linked to these classifications are much less likely to receive targeted attacks than would be expected.

DISCUSSION

Targeted attacks tend to be mentioned in slightly hysterical press headlines relating to the latest large organization to succumb, or discussed as part of the development of a new and sinister cyber cold war. It is clear that high-value information and systems present tempting targets to sophisticated attackers, but there is nothing inevitable about such attacks being successful. Protection against such threats requires the development of robust systems that are, as much as possible, resistant to attack, coupled with constant monitoring for evidence of intrusion, and mitigation strategies to resolve the attack when detected.

Advance identification of where attacks are mostly likely to be directed allows a higher degree of protection to be allocated where needed. High-risk individuals and systems may require a tailored information security programme appropriate to their needs. This may take the form of different security policies, additional security software, and enhanced training, allowing budgets to be concentrated where they may be most effective. However, to achieve this, the techniques by which these high-risk individuals can be identified must be developed.

The methodology of case control studies [11] can be adopted from traditional epidemiology and applied to information security to investigate putative risk factors for outcomes such as being sent targeted attacks. In such studies we can search for putative factors that are associated at more than 95% confidence with an outcome. The tentative identification of such factors can be used to design more powerful epidemiological studies, such as cohort or randomized control studies to better calculate relative risk [16, 17].

Much work on the epidemiology of malware has considered the spread of self-replicating malware across vulnerable systems [18–20]. However, this is to ignore the effect of trojan malware that does not attempt to propagate, and does not consider the effect of the traits associated with the individual whose system becomes infected. Carlinet *et al.* have used epidemiological techniques to identify risk factors for ADSL users to generate malicious traffic. The study identified that the use of web and streaming applications and use of the *Windows* operating system were risk factors for apparent malware infection [21]. Bossler and Holt conducted a similar study looking at factors associated with malware infection, finding that media piracy was positively

correlated with infection, as was '*associating with friends who view online pornography*', being employed and being female [22].

In these cases it appears as if user behaviour is leading to increased exposure to malware and increasing risk of infection. Targeted trojans differ from other common forms of malware in that the attacker researches and selects potential targets to which the attacks are directed. It is not necessarily the behaviour of the individual that leads to exposure to malware, but rather that something specific to the individual leads them to come to the attention of attackers who then launch attacks against the target.

A reasonable hypothesis is that it is an individual's area of expertise that leads to them becoming of interest to attackers and becoming subject to targeted attacks. In this case, it should be possible to discover the predilections of the attackers for certain subjects over others in analysing the differences in the profile of recipients of non-targeted and targeted malware.

Indeed, the data presented supports this hypothesis. For recipients of malware in the academic sector, individuals working in Eastern, Asiatic, African, American and Australasian Languages, Literature and Related Subjects and Social Studies, especially Economics, are at a statistically significant increased risk of being subjected to targeted attacks, odds ratios 12.03 and 11.79 respectively.

This is not to say that individuals working in these subjects will be subject to targeted attacks, but that they are at an increased risk. Equally, we cannot infer that working in these subjects has directly caused the individual to become targeted, merely that there is an association. However, the discovery of associations is often the first step in discovering the steps involved in causation.

Identification of the further features that contribute to an individual being subject to attack requires understanding not only the factors within the target population that predispose them to attack, but also understanding the factors within the population of attackers that cause them to attack some individuals rather than others.

CONCLUSION

Targeted attacks are amenable to study as a public health issue by comparing the likelihood that features are found in the set of recipients of such attacks to recipients of non-targeted attacks. Achieving this requires the selection of a suitable control population, a selection of features to test, and a suitable methodology.

Applying a case control study to academic malware recipients, using the HESA JACS coding of academic subjects to investigate the relationship between research interests and the receipt of targeted attacks, shows that there is statistically significant correlation for being subject to targeted attacks for researchers in Eastern, Asiatic, African, American and Australasian Languages, Literature and Related Subjects and Social Studies. Conversely, researching Medicine & Dentistry, Veterinary Science, Agriculture and Related Subjects, Architecture Building & Planning, and being a non-academic staff member are apparent protective factors for being subject to targeted attacks.

Considering traits within populations as potential risk or protective factors for attack can lead to the identification of individuals who are at high risk of being subject to sophisticated attacks. In turn, this may allow the concentration of defences where attacks are most likely to occur and result in increased protection.

ACKNOWLEDGEMENTS

This paper would not have been possible without the hard work of Tony Millington. Thanks to Steve White, Paul Wood and Alistair Johnson for their continued support, and to Olivier Thonnard for his insightful feedback.

REFERENCES

[1] Symantec, Information Security Threat Report. Vol. 17 (2012). http://www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_2011_21239364.en-us.pdf.

[2] MacSweeney, G. The Top 9 Most Costly Financial Services Data Breaches. Wall Street and Technology (April 2012). http://www.wallstreetandtech.com/data-security/232800079.

[3] Schwartz, M.J. RSA SecurID Breach Cost $66 Million. Information Week (July 2011). http://www.informationweek.com/news/security/attacks/231002833.

[4] Ross, R.; Katzke, S.; Johnson, A.; Swanson, M.; Stoneburner, M.; Stoneburner, G. Managing risk from information systems: An organizational perspective. NIST Special Publication 800-39 Appendix B. 2011. http://csrc.nist.gov/publications/nistpubs/800-39/SP800-39-final.pdf.

[5] Epidemiology. Merriam-Webster.com (May 2012). http://www.merriam-webster.com/epidemiology.

[6] Newsom, S.W.B. Pioneers in infection control: John Snow, Henry Whitehead, the Broad Street pump, and the beginnings of geographical epidemiology. Journal of Hospital Infection, Vol. 64(3) (November 2006), pp.210–216, ISSN 0195-6701, DOI: 10.1016/j.jhin.2006.05.020. http://www.sciencedirect.com/science/article/pii/S0195670106002830.

[7] Doll, R.; Hill, A.B. Lung cancer and other causes of death in relation to smoking; a second report on the mortality of British doctors. British Medical Journal Vol. 2(5001) (November 1956), pp.1071–1081. DOI: 10.1136/bmj.2.5001.1071. PMC: 2035864. PMID: 13364389. http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2035864/.

[8] Doll, R.; Peto, R.; Wheatley, K.; Gray, R.; Sutherland, I. Mortality in relation to smoking: 40 years' observations on male British doctors. British Medical Journal Vol. 309(6959) (October 1994) pp.901–11. PMCID: PMC2541142. http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2541142/.

[9] Targeted Trojan Email Attacks. Centre for the Protection of National Infrastructure Briefing 08/2005. (June 2005). http://www.cpni.gov.uk/Documents/Publications/2005/2005015-BN0805_Targeted_trojan_email.pdf.

[10] Bland, J.M.; Altman, D.B. Statistics Notes: The odds ratio. British Medical Journal. Vol. 320(7247) (May 2000) p.1468. PMCID: PMC1127651. http://www.ncbi.nlm.nih.gov/pmc/articles/PMC1127651/.

[11] Schultz, K.F.; Grimes, D.A. Case-control studies: research in reverse. The Lancet, Vol. 359(9304) (February 2002) pp.431–34. PMID: 11844534. https://research.chm.msu.edu/Resources/5%20case%20control%20studies.pdf.

[12] Morris, J.A.; Gardner, M.J. Statistics in Medicine: Calculating confidence intervals for relative risks (odds ratios) and standardised ratios and rates. British Medical Journal, Vol. 296(6632) (May 1988) pp.1313–1316. PMCID: PMC2545775. http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2545775/.

[13] Lee, M.; Lewis, D. Clustering Disparate Attacks: Mapping the Activities of the Advanced Persistent Threat. Proceedings of the 21st Virus Bulletin International Conference. (October 2011) pp.122–127.

[14] Thonnard, O.; Bilge, L.; O'Gorman, G.; Kiernan, S.; Lee, M. Industrial Espionage and Targeted Attacks: Understanding the Characteristics of an Escalating Threat. To appear in Proceedings of 15th International Symposium on Research in Attacks Intrusions and Defenses (September 2012).

[15] Full JACS3 Listing v.1.2. Higher Education Statistics Agency (September 2011). http://www.hesa.ac.uk/content/view/1805/277/.

[16] Szklo, M. Population-based Cohort Studies. Epidemiologic Reviews. Vol. 20(1) (1998) pp.81–90. http://epirev.oxfordjournals.org/content/20/1/81.full.pdf.

[17] Altman, D.G.; Bland, J.N. Statistics Notes: Treatment allocation in controlled trials: why randomise? British Medical Journal, Vol.318(7192) (May 1999) p.1209. DOI: 10.1136/bmj.318.7192.1209. PMID: 10221955. http://www.bmj.com/content/318/7192/1209.1.

[18] Kephart, J.O.; White, S.R.; Chess, D.M. Computers and Epidemiology. Spectrum, IEEE. Vol. 30(5) (1993) pp.20–26. DOI: 10.1109/6.275061. http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=275061.

[19] Martin, J.C.; Burge, L.L. III; Gill, J.I.; Washington, A.N.; Alfred, M. Modelling the spread of mobile malware. International Journal of Computer Aided Engineering and Technology. Vol.2 (2010) pp.3–14. DOI: 10.1504/IJCAET.2010.029592. http://www.inderscience.com/search/index.php?action=record&rec_id=29592.

[20] Garetto, M.; Gong, W.; Towsley, D. Modeling malware spreading dynamics. INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies Vol. 3 pp.1869 – 1879. DOI: 10.1109/INFCOM.2003.1209209. http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=1209209.

[21] Carlinet, Y.; Me, L.; Debar, H.; Gourhant, Y. Analysis of Computer Infection Risk Factors Based on Customer Network Usage. Emerging Security Information, Systems and Technologies, 2008. SECURWARE '08. pp.317–325, (August 2008) DOI: 10.1109/SECURWARE.2008.30. http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=4622601.

[22] Bossler, A.M.; Holt, T.J. On-line Activities, Guardianship, and Malware Infection: An Examination of Routine Activities Theory. International Journal of Cyber Criminology. Vol. 3(1) (2009) pp.400–420. http://www.cybercrimejournal.com/bosslerholtijcc2009.pdf.