Phishing and Organisational Learning

WD Kearney, HA Kruger
School of Computer, Statistical and Mathematical Sciences,
North-West University, Private Bag X6001, Potchefstroom, 2520
South Africa
Kearneys@iinet.net.au, Hennie.Kruger@nwu.ac.za

Abstract. The importance of addressing the human aspect in information security has grown over the past few years. One of the most frequent techniques used to obtain private or confidential information from humans is phishing. One way to combat these phishing scams is to have proper security awareness programs in place. In order to enhance the awareness and educational value of information security awareness programs, it is suggested that an organisational learning model, characterised by so called single-loop and double-loop learning, be considered. This paper describes a practical phishing experiment that was conducted at a large organisation and shows how a learning process was initiated and how security incidents such as phishing can be used successfully for both single and double-loop learning.

Keywords: Phishing, Social engineering, Information security awareness, Organisational learning.

1 Introduction

Traditionally the mitigation of information security risks was addressed using a variety of technical controls. It is however widely accepted and recognised that technology on its own cannot deliver complete solutions to the security problem and that the human aspect of security should receive more attention [1], [2], [3]. One way of addressing the human side of security is to focus on awareness and educational activities [4] making use of some form of an awareness program.

An information security awareness program normally focuses on a number of issues related to the correct security behaviour of users. In some instances it may also concentrate on one area such as social engineering which is one of the most serious threats to information security as criminals keep on focussing on deceptive techniques to attack computer users and organisations [5]. Phishing, which is one of the social engineering techniques, occurs when people are manipulated by deception into giving out information [6] and is one of the major threats to modern organisations and information technology users in general. It requires an ongoing awareness not to become a victim of a phishing scam and various researchers have completed studies related to phishing experiments and awareness levels of users [5], [7], [8].

A popular technique to improve user awareness pertaining to phishing scams is to conduct unannounced phishing tests in order to evaluate users' propensity to respond to an attack [5], [9]. Albrechtsen [10] contend that these type of incidents and experiments present great opportunities to learn and improve information security. To ensure that learning does take place Van Niekerk and Von Solms [3] suggested that an organisational learning model be used.

This paper describes a practical phishing exercise that was conducted in industry and shows how organisational learning took place as a result. The remainder of the paper is organised as follows. Section 2 presents the background to the study as well as appropriate references to related work. In section 3 the methodology used is discussed while section 4 details the results. Concluding remarks are presented in section 5.

2 Background and related work

Organisational learning theories deal with the idea of how organizations learn and adapting its behaviour [3]. This concept has been subjected to a wide and growing variety of researchers and a number of definitions have been suggested in the literature [11], [12]. Despite all these definitions the concept of organizational learning is by no means an unambiguous concept, as no one irrefutable definition has emerged in literature [13]. Organisational learning originated from the work by Argyris and Schon during the 1970s and one of the definitions suggested by them will be assumed in this study. The definition is formulated as follows. Organisational learning occurs when individuals within an organisation experience a problematic situation and enquire into it on the organisational behalf [14].

In an effort to enhance organisational learning, Buckler [15] proposed that an actual learning process, as depicted in figure 1, occurs in organisations. Buckler then argues that individuals will move through the different learning stages driven by their inherent individual motivations to learn. Associated with these motivational forces, there will be certain barriers to the learning process, and where the motivational restraining (barrier) forces are matched, learning will not take place. In order for organisational learning to result in performance improvement, the enactment stage (see figure 1) of the learning process needs to be achieved – this will imply behavioural change which is a requirement for successful organisational

learning. To assess the effectiveness of the behavioural changes, the reflection stage should be entered.

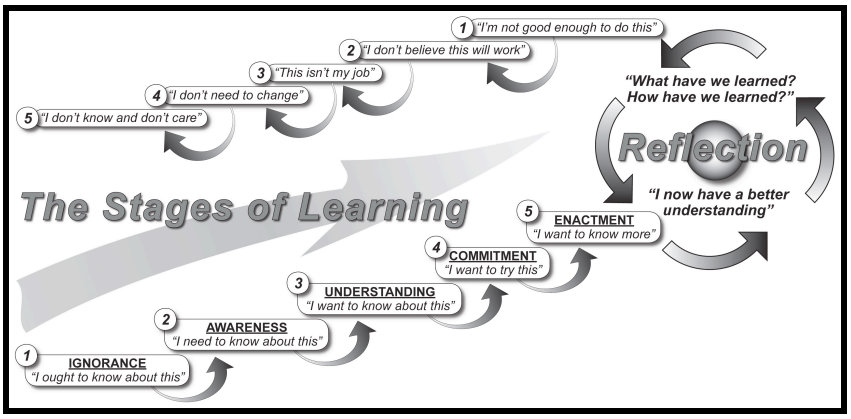

Fig. 1. The learning process (adapted from [15])

There are various applications of learning processes but in general three types of learning can be categorised. These three types are summarized by Kennedy [13] as follows.

- Single-loop learning, which occurs when errors are detected and corrected and organisations continue with the present status quo without modifying present policies and goals. In essence, single-loop learning focuses on improving the status quo through small incremental changes in how organisations functions. An example in the area of information security could be a case of unauthorized access by a user to privileged data. A single-loop response would be to simply deny future access to this specific user. The status quo is maintained and present policies and/or goals are not modified.
- Double-loop learning challenges, and possibly makes changes to the status quo and the existing assumptions and conditions. It means that the organisation questions and modifies its existing norms, policies, procedures and objectives and it can lead to transformational change that radically alters the status quo. In the information security example mentioned, a double-loop response may be to investigate the circumstances and reasons for the unauthorized access. Double-loop learning may then occur when a decision is taken to improve (change) the process of allocating access rights in order to minimize future unauthorized access risks.
- Deutero learning involves focusing on the learning process itself. This type of learning seeks to improve how organisations perform single and double-loop learning. It can be described as "learning how to learn" and it occurs when organisations learn how to perform both single and double-loop learning.

Due to the focus on long term goals and the more complex nature of double-loop learning, most companies focus only on single-loop learning [16]. According to Van Niekerk and Von Solms [3] this is also true in the information security discipline.

They pointed out that generative, or double-loop learning, emphasizes continuous experimentation and feedback.

Although there are a large number of studies on organisational learning, there are not particularly many studies that relate organisational learning to information security. Even so, the studies that have been conducted in this area prove that information security is an important area that offered ample opportunities, linked to organisational learning, that can make a significant contribution to organisations and their performance. Examples of studies where organisational learning and information security were explored include the following.

Van Niekerk and Von Solms [3] investigated, amongst other models, the use of an organisational learning model for information security education. Their aim was to ensure that adequate attention is given to behavioural theories in information security education programs. Albrechtsen [10] conducted a comprehensive study into the barriers that exist and that prohibit productive organisational learning from information security incidents while Ahmat *et al* [16] suggested that the practice of incident response may lead to organisational learning. They proposed a double-loop learning model for security incident learning to address potential systemic corrective action. An interesting and authoritative study was conducted by Pfleeger and Caputo [2] where it was argued that blending behavioural sciences and cyber security may lead to the mitigation of cyber security risks. Although organisational learning was not specifically mentioned, the study strongly supports the idea that behavioural sciences (of which organisational learning is at least a sub-section) is relevant to information security in general. To further motivate this idea, Thomson and Van Niekerk [4] also contend that employee apathy towards information security can be addressed through the use of existing theory from the social sciences.

There are also a number of studies where the focus is not on information security per se but rather on how information technology in general relates to organisational learning. These studies usually concentrate on computer systems necessary to facilitate organisational learning and knowledge transfer [17], [18].

In the context of this paper, where it is claimed that a phishing exercise may lead to organisational learning, the next few paragraphs will briefly refer to the phishing concept and examples of studies related to it.

The basic idea of phishing is when someone attempts to fraudulently acquire sensitive information from a victim by impersonating a trustworthy entity [8]. A more formal definition can be obtained from the Oxford English Dictionary [19] where phishing is defined as *the fraudulent practice of sending e-mails purporting to be from reputable companies in order to induce individuals to reveal personal information, such as passwords and credit card numbers, online.*

Phishing attacks are on the increase and successful attacks may have devastating effects on both enterprises and individuals. The Symantec Intelligence Report [20] of June 2012 reported that one out of every 170.9 e-mails sent during the month of June 2012, in South Africa, was a phishing scam. In the Netherlands the figure for June 2012 was one out of every 54.4 e-mails. Considering the billions of e-mail messages that are transmitted worldwide during a specific month, it becomes clear to what extend phishing attacks form part of the day to day electronic communication activities. With this in mind it becomes more and more important to implement the right and effective countermeasures to mitigate or prevent phishing attacks. One way

of dealing with this growing number of phishing incidents is to implement security awareness and training programs where users are made aware of phishing scams. The use of practical tests seems to be a popular and effective way of making people aware of the dangers of phishing and some examples of the work conducted by other researchers in this area will be highlighted below.

Pattison *et al* [21] investigated the behaviour response of computer users when receiving either phishing e-mails or genuine e-mails. The study was conducted as a scenario-based role-play experiment where participants had to indicate what the appropriate response would be on certain e-mail messages. The study found that participants who were informed, prior to the experiment, that they are part of a phishing exercise performed better in handling phishing e-mail messages.

Simulated phishing attacks together with embedded training were used by Jansson and Von Solms [5] in an effort to cultivate users' resistance towards phishing attacks, while Kumaraguru *et al* [7] also conducted a study on anti-phishing training to proof that user training should be used in conjunction with technological solutions for security problems. Other studies include Dodge *et al* [9] who performed a practical phishing experiment involving students from the United States Military Academy, Jagatic *et al* [8] performed a study at the Indiana University, Steyn *et al* [22] conducted a practical experiment in South Africa and Hasle *et al* [23] a study in Norway.

It is interesting to note that all the practical phishing experiments referred to so far, were conducted using students as participants. Although these studies produced many advantages and insights, it is doubted whether the results can be generalised and extrapolated to industry enterprises.

Consistent with the research projects mentioned above, this study also performs a practical phishing experiment but uses an industry enterprise for research purposes instead of students in a university environment. In addition, the exercise is aimed at creating a climate for organisational learning. To ensure that the exercise is not a once-off event, the objective is to initiate a learning process and to show how security incidents such as phishing can and should be used for single and double-loop learning in an organisation.

The study was conducted at a large geographically dispersed utility. The organisation in question is a large multi-billion dollar entity with over 3500 IT users and they supply essential services to over 2 million customers. The organisation has an information security course that is mandatory for all employees and partners who have access to the IT infrastructure. The objective of the course is to make IT users aware of their responsibilities with regards to protecting the organisations' information and information systems from unauthorised access, loss or disclosure. Whilst the information security course is deemed mandatory, the records could not support this assertion as many staff was found not to have completed the course or no records could be found of their attendance.

3 Methodology

The successful implementation of an e-mail phishing exercise is dependent on how well certain issues, associated with the exercise, are considered. Jansson and Von

Solms [5] categorised these issues into principles to be considered *before designing* the exercise, *before conducting* the exercise, *during* the exercise, and *after* the exercise while Dodge *et al* [9] simply refer to them as general and specific considerations. In this study considerations are also presented as general and specific considerations. The general considerations are concerned with those issues that may have an impact on the exercise as a whole while the specific considerations deal with aspects specific to the enterprise where the study was conducted.

General considerations

The first and most important general consideration is the determination and definition of an objective. There should be a clearly defined goal and in this study the goal was simply stated as the evaluation of security awareness associated with phishing and the creation of an opportunity for organisational learning to take place. The next consideration is critical for success i.e. to get ethical clearance and top management approval. This was achieved by conducting personal meetings with the CEO, the CFO and the IT manager where the purpose, actual steps and possible outcomes were explained. A formal project proposal detailing aspects such as the basic process, different phases, measures of success and possible risks, was also submitted for approval to management.

Other general considerations which were appropriately addressed included the timing of the exercise; maintaining the privacy of respondents; the selection of a random and representative sample of respondents; measurements to ensure that no information was disclosed prior to the exercise; and, a debriefing exercise following the test.

Specific considerations

The central issue among the specific considerations was the construction of an appropriate e-mail message. The message had to be concise, credible and at the same time be enticing in order for participants to react.

To ensure that the phishing e-mail message complies with all the necessary requirements, it was decided to make use of aspects that may trigger certain emotions from participants. Jansson [6] presents a list of a large number of techniques that are based on negative, positive and neutral emotional exploits. For the construction of the e-mail message the following emotional exploits were used.

Legitimacy – when a user is made to believe that the source of the e-mail message is legitimate.

Authority – people tend to comply with instructions or requests issued by someone with authority.

Scarcity – when users believe that the time to react is limited.

Conformity – users who believe that other fellow-employees have already reacted to a request are inclined to also comply with the request.

Apart from these four techniques which were explicitly built into the e-mail message (see figure 2), three other important emotional exploits were also implicitly included. They were *urgency* (making users believe it is an emergency), *carelessness* (clicking on a link) and *diffusion of responsibility* (users believe that someone else is responsible for security). Users were asked to click on the link in the message which would then take them to another webpage where their usernames and passwords

were requested. Figure 2 also indicates how the e-mail was constructed to provide clues to alert users that the message was likely not to be legitimate. The real name of the organisation has been changed in figure 2.

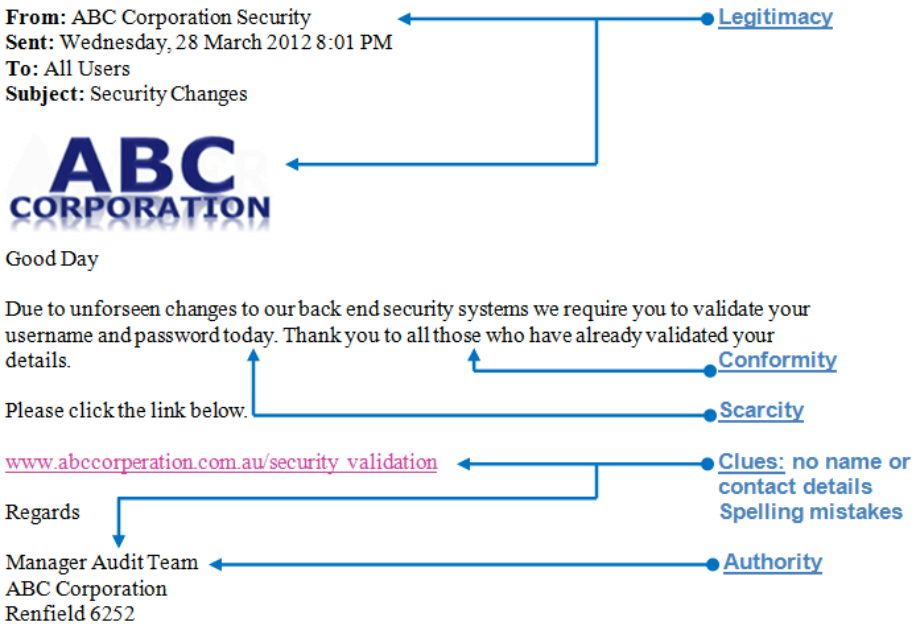

Fig. 2. Phishing e-mail message

There were a number of other specific issues that also needed clarification before the actual exercise could take place i.e. it was important not to refer to any specific IT, security or internal audit staff as this may compromise the trust between users and staff. Steps also had to be taken to ensure that the enterprise's anti-phishing tools and spam filters do not identify the message as spam or a phishing scam, and Helpdesk had to be provided with a predetermined response should there be any queries from users. Provision was also made for respondents who reply directly to the phishing e-mail. Some of the technical considerations include the deletion of duplicate records (if a user responds more than once) and also a check to see whether the correct usernames were supplied (password were requested but not recorded).

The e-mail message (figure 2) was first sent to a small group of 10 employees. The objective was to test whether all technical aspects are functioning correctly and also to get feedback on possible improvements. After some minor changes were made, following the small pilot study, it was decided to go ahead and implement the phishing test.

The phishing e-mail message was sent to all employees at 8:00pm on a weekday night. The organisation is a 24-hour operation with activities taking place on a continuous basis. Statistics of user logs showed that there are on average about 1700 active IT users signed on during any night and to ensure that the night workers are included in the test, the 8:00pm sending time was chosen. This sending time would also guarantee that day workers should have the phishing e-mail in their inboxes first

thing in the morning. The idea was to get users to respond early before they can discuss it with fellow employees.

A number of senior managers found the phishing e-mail very annoying and some of them sent out general e-mail messages to object to the phishing message (and the test). The security personnel were also involved and concern was expressed regarding the possibility of an external attack aimed at disrupting essential services. Due to this, it was decided at 8:30am the next morning to remove the phishing message and to officially end the test. The reasons for withdrawing the phishing e-mail relatively early the next morning were firstly, to prevent large-scale disruptions and secondly, because enough data has been recorded at that stage to draw meaningful conclusions. The data and the experience were sufficient and interesting results, presented in the next section, were obtained.

4 Results

The data recorded from the phishing awareness exercise include the employee name, department where the person is working and the username. Passwords were also requested but not recorded due to privacy considerations. As part of the exercise, passwords were validated but only the result was recorded in a simple yes/no format. Appropriate safeguards to ensure privacy were put in place. The recorded employee names were purely recorded for statistical purposes and nowhere during reporting were specific names linked to responses. The reason for recording usernames was to perform a validation test to ensure that users do enter valid usernames (and by implication valid passwords). All duplicate records (users who entered their details more than once) and records with invalid usernames were removed from the final data set.

The main result, before any further analyses were performed, was the number of negative responses received. A negative response is a response where a user provided his or her username and password. During the test 280 users responded to the phishing message of whom 231 (83%) entered their usernames and passwords on the webpage. Of the 231 users, 23 (10%) entered their valid details more than once. Although there were approximately 1700 active users logged on during the test, it would be incorrect to assume that all of those who did not respond acted in a positive way. Reasons for this may be the fact that many people do not respond immediately to e-mail messages, some users may have left their workstations logged on during the night while not there, some users may have been engaged in other tasks and simply did not check their mail inboxes, etc. A much more significant analysis was to link the 280 users who responded, to the information security course that all staff members are required to complete and which would have provided them with basic security information on how to react to possible phishing scams. Figures 3(a) and (b) show the results graphically. Figure 3(a) shows that an unexpected 69% of those users who entered their passwords did complete the security training in the past. Figure 3(b) shows the training details of those who responded without entering their usernames and passwords. These results indicate that there are at least two points of concern. Firstly, the high number of users who responded in a negative way despite

their security training and secondly, the relatively high number of users that never completed the information security course.

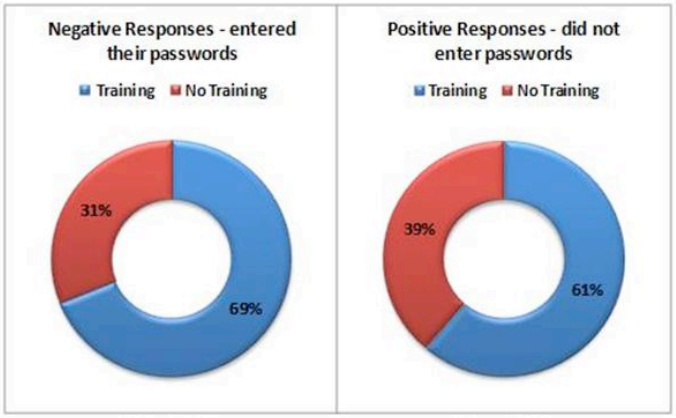

Fig. 3. Responses related to training completed

Figure 4 shows an analysis of responses (percentages) per experience category for those who entered their usernames and passwords. Experience in this case refers to the number of years a person is employed at the organisation. From figure 4 it can be seen that those employees with less experience at the organisation (and therefore less exposure to its security practices and policies) are more inclined to give away personal details. More than a third (35%) of those who entered their usernames and passwords have less than 5 years experience with more than half (52%) less than 10 years.

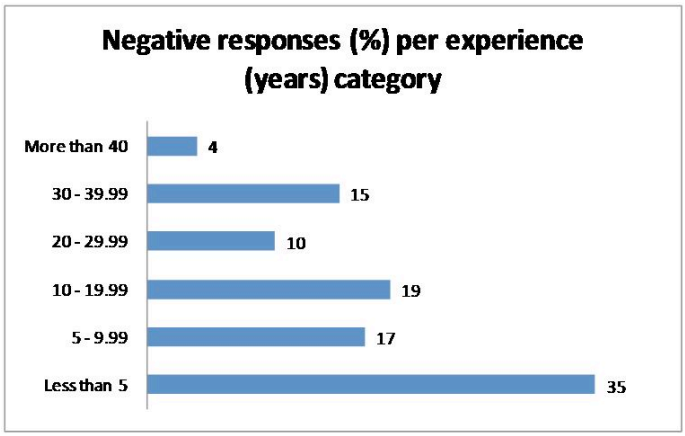

Fig. 4. Responses per experience category

The data that was captured during the exercise makes it possible to perform a number of analyses, e.g. responses per department, gender, age group etc. These types of analyses were not done in this study as the focus was more directed at possible organisational learning opportunities.

As explained earlier, organisational learning involves the adjustment of actions based on an experience. These adjustments, or learning, can then be categorised as single or double-loop learning. The results from this study have shown that the phishing experiment offers the ideal opportunity for learning and that both single and double-loop learning has taken place.

Single-loop learning took place in the form of small changes in making staff aware of the risks and consequences of phishing scams. Instructions concerning basic acceptable behaviour related to suspicious e-mail messages were also issued. Specific actions that can be attributed to single-loop learning include the following.

- The first day, following the phishing exercise, the Manager Risk and Assurance sent out an e-mail message to all staff informing them about the exercise and, more importantly, making them aware of the risks and giving them basic instructions on how to react to these type of e-mails (e.g. to report it to the Service Centre).
- The company's weekly in-house bulletin was used to reinforce the security awareness message and to instruct staff to complete the company's computer based information security course. This was done for two consecutive months following the phishing exercise.

The single-loop learning examples mentioned here did not change the status quo of any process but were quick and effective corrective measures to address a specific problem area. There were, however, other issues that needed a more comprehensive investigation that may lead to a change in policies and procedures. These double-loop learning issues include the following.

- All staff members are required to complete an information security course which will equip them with basic security knowledge for different security situations including phishing scams. An analysis of the phishing results showed that not all staff has completed the course. More importantly, a relatively large number of those who have completed the course had given their passwords away. An assessment of the course content and possible controls to ensure that everybody completes the course is planned. This may lead to a change in the current security policy on issues pertaining to basic security training.
- Another issue, planned for the future, which was highlighted during the phishing exercise relates to the gap between the different security views and expectations of managers and users. This gap is sometimes referred to as the information security digital divide between managers and users [24] and may lead to unrealistic security assumptions and management strategies that are not aligned with the dynamics of the user environment.

If one considers the results of the phishing exercise it seems permissible to draw the conclusion that the exercise has created opportunities for organisational learning. Basic problems were immediately corrected through an easy and uncomplicated

single-loop learning approach while double-loop learning issues provided an opportunity for the organisation to adapt and adjust some of their information strategies.

5 Conclusions

Modern businesses are characterised by the increasing reliance on information assets. The protection of these assets depends to a large extend on the employees and users and it is not surprisingly that criminals tend to focus their attacks on humans. Phishing has become one of the most frequently used techniques to obtain personal or private information and to combat it, proper security awareness programs should be in place. To ensure that a security awareness activity does not become a once-off event, organisations may want to consider the use of various organisational learning models to enhance the awareness and educational value of such programs.

In this paper a successful practical phishing exercise was conducted at a large organisation. The aim was not only to record the number of users who are willing to give away personal information, but also to create an opportunity for organisational learning in order to improve the educational value of the phishing experiment. The results have shown that employees are prone to phishing attacks, but more importantly, the phishing exercise created an excellent opportunity for both single and double-loop learning activities. A single-loop learning approach was followed to immediately correct certain shortcomings without changing the status quo, while double-loop learning provided the opportunity to revisit and adapt some of the longer term information security strategies.

One security experiment linked successfully to organisational learning does not necessarily prove that all security exercises will lead to organisational learning. The exercise did, however, provide an insight into exciting possibilities to increase the value of security awareness exercises and that it may ultimately lead to the completion of the learning process described in section 2 of the paper.

References

1. Furnell, S., Clarke, N.: Power to the People? The Evolving Recognition of Human Aspects of Security, Computers and Security 31, 983-988 (2012)
2. Pfleeger, S.L., Caputo, D.D.: Leverage Behavioral Science to Mitigate Cyber Security Risk, Computers and Security 31, 597-611 (2012)
3. van Niekerk, J., von Solms R.: Organisational Learning Models for Information Security, (2004); http://icsa.cs.up.za/issa/2004/Proceedings/Full/043.pdf.
4. Thomson, K., van Niekerk, J.: Combating Information Security Apathy be Encouraging Prosocial Organisational Behaviour, Information Management & Computer Security 20, 39-46, (2012)
5. Jansson, K., von Solms, R.: Phishing for Phishing Awareness, Behaviour & Information Technology, DOI:10.1080/0144929X.2011.632650, (2011)
6. Jansson, K.: A Model for Cultivating Resistance to Social Engineering Attacks, Unpublished M-dissertation, Nelson Mandela Metropolitan University, (2011)

7. Kumaraguru, P., Cranshaw, J., Acquisti, A., Cranor, L., Hong, J., Blair, M.A., Pham, T.: School of Phish: A Real-World Evaluation of Anti-Phishing Training, In: Proceedings of the 5th Symposium on Usable Privacy and Security (SOUPS), pp. 3:1-3:12, (2009)
8. Jagatic, T.N., Johnson, N.A., Jakobsson, M., Menezer, F.: Social Phishing, Communications of the ACM, 50(10):94-100 (2007)
9. Dodge, R.C., Carver, C., Ferguson, A.J.: Phishing for User Security Awareness, Computers and Security 26,73-80 (2007)
10. Albrechtsen, E.: Barriers against Productive Organisational Learning from Information Security Incidents, Paper in the PhD course Organisational Development and ICT, Norwegian University of Science and Technology, (2003)
11. Schermerhorn, J.R., Osborn, R.N., Uhl-Bien, M., Hunt, J.G.: Organisational Behavior, 12th edition, (John Wiley & Sons, Inc, NJ, 2012)
12. Lopez, S.P., Peon, J.M.M., Ordas, C.J.V.: Organisational Learning as a Determining Factor in Business Performance, The Learning Organisation 12(3), 227-245 (2005)
13. Kennedy, E.: A Critical Evaluation of the Organisational Learning that takes place in a Project Management Environment, Unpublished M-dissertation, North-West University (2008)
14. Argyris, C., Schon, D.: Organisational Learning II: Theory, Method and Practice, (Prentice Hall, 1996)
15. Buckler, B.: Practical Steps towards a Learning Organisation: Applying Academic Knowledge to Improvement and Innovation in Business Processes, The Learning Organisation 5(1),15-23, (1998)
16. Ahmad, A., Hadgkiss, J., Ruighaver, A.B.: Incident Response Teams – Challenges in Supporting the Organisational Security Function, Computers and Security 31, 643-652 (2012)
17. Kane, G.C., Alavi, M.: Information Technology and Organisational Learning: Investigation of Exploration and Exploitation Processes, Organization Science 18(5),796-812, (2007)
18. Chou, S.: Computer Systems to Facilitating Organizational Learning: IT and Organizational Context, Expert Systems with Applications (24), 273-280 (2003)
19. Oxford Dictionary (November 2012), http://oxforddictionaries.com/definition/english/phishing
20. Symantec Intelligence Report (June 2012), www.symantec.com/content/en/us/enterprise/other_resources/b_intelligence_report_06_2012.en-us.pdf
21. Pattinson, M., Jerram, C., Parsons, K., McCormac, A., Butavicius, M.: Why do some People Manage Phishing E-mails Better than Others?, Information Management and Computer Security 20(1), 18-28 (2012)
22. Steyn, T., Kruger, H.A., Drevin, L.: Identity Theft – Empirical Evidence from a Phishing Exercise, New Approaches for Security, Privacy and Trust in Complex Environments, IFIP International Federation for Information Processing (232), 193-203 (2007)
23. Hasle, H., Kristiansen, Y., Kintel, K., Snekkenes, E.: Measuring Resistance to Social Engineering, In: ISPEC05 Proceedings of the First International Conference on Information Security Practice and Experience, 132-143 (2005)
24. Albrechtsen, E., Hovden, J.: The Information Security Digital Divide between Information Security Managers and Users, Computers and Security (28), 476-490 (2009)