COMPUTERS & SECURITY 69 (2017) 18–34

Available online at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/cose

Computers & Security

Panning for gold: Automatically analysing online social engineering attack surfaces

Matthew Edwards *, *Robert Larson*, *Benjamin Green*, *Awais Rashid*, *Alistair Baron* *

Security Lancaster, School of Computing and Communications, Lancaster University, Lancaster LA1 4WA, UK

ARTICLE INFO

Article history:
Available online 29 December 2016

Keywords:
Social engineering
Vulnerability analysis
Open source intelligence
Social networks
Competitive intelligence

ABSTRACT

The process of social engineering targets people rather than IT infrastructure. Attackers use deceptive ploys to create compelling behavioural and cosmetic hooks, which in turn lead a target to disclose sensitive information or to interact with a malicious payload. The creation of such hooks requires background information on targets. Individuals are increasingly releasing information about themselves online, particularly on social networks. Though existing research has demonstrated the social engineering risks posed by such open source intelligence, this has been accomplished either through resource-intensive manual analysis or via interactive information harvesting techniques. As manual analysis of large-scale online information is impractical, and interactive methods risk alerting the target, alternatives are desirable.

In this paper, we demonstrate that key information pertinent to social engineering attacks on organisations can be passively harvested on a large-scale in an automated fashion. We address two key problems. We demonstrate that it is possible to automatically identify employees of an organisation using only information which is visible to a remote attacker as a member of the public. Secondly, we show that, once identified, employee profiles can be linked across multiple online social networks to harvest additional information pertinent to successful social engineering attacks. We further demonstrate our approach through analysis of the *social engineering attack surface* of real critical infrastructure organisations. Based on our analysis we propose a set of countermeasures including an automated social engineering vulnerability scanner that organisations can use to analyse their exposure to potential social engineering attacks arising from open source intelligence.

1. Introduction

Social engineering attacks pose a major risk to the security of organisations. Some of the most high profile cyber attacks on large organisations, e.g., RSA, JP Morgan, AT&T, the Ukranian power grid, etc., leveraged social engineering as an entry point into the organisation's systems. Attackers use a number of tactics, ranging from simple impersonation to complex multi-layered deceptions worthy of a Hollywood caper, that lead a target to disclose sensitive information or to interact with a malicious payload. At their most basic, these attacks may be represented by a generic phishing email from an unfamiliar sender that targets hundreds of staff within an organisation

* Corresponding authors.
E-mail addresses: m.edwards7@lancaster.ac.uk (M. Edwards), r.larson@lancaster.ac.uk (R. Larson), b.green2@lancaster.ac.uk (B. Green), a.rashid@lancaster.ac.uk (A. Rashid), a.baron@lancaster.ac.uk (A. Baron).
http://dx.doi.org/10.1016/j.cose.2016.12.013

with the same message. More sophisticated attacks may greatly increase their chance of success by targeting a much smaller pool of recipients with a personalised ploy (Jagatic et al., 2007).

Current research suggests that the effectiveness of such attacks can be greatly increased through the use of open source intelligence (OSINT) to boost the effectiveness of the deceptive ploys delivered in an attack (Jagatic et al., 2007). Such open source information is now widely available – with individuals increasingly releasing information about themselves online, particularly on social networks. Even more worryingly, practices such as organisational engagement with social media and the publication of employee rosters on organisational websites are enabling attackers to easily identify an organisation's employees from amongst millions of social media users. This lets attackers know exactly who to target for data harvesting in preparation for an attack on the organisation. Methods by which such OSINT data may be used to increase effectiveness in this manner include (but are not limited to): selection of vulnerable personalities, inclusion of ploys personally attractive to the target, and impersonation of a person in authority (Huber et al., 2009).

Existing research has demonstrated the social engineering risks posed by such OSINT data (Ball et al., 2012). However, this normally relies on labour intensive manual analysis (Creese et al., 2012), which is impractical and poses a high cost to a potential attacker. Alternatively, such techniques utilise automated conversational agents (Huber et al., 2009), which do not scale and are not very effective due to the challenges of imitating human conversational behaviour. Other techniques rely on "active" engagement with potential targets to elicit information – through zombie profiles or misleading friend requests (Scheelen et al., 2012) – and hence risk detection prior to an attack being launched. In this paper, we demonstrate that both of these challenges – automation and passive information gathering – can be overcome, posing major social engineering risks to organisations.

We show that it is possible to automatically identify the employees of an organisation amongst individuals within its online footprint. Furthermore, we demonstrate that it is possible to automatically resolve employee identities across multiple online social networks, with a high accuracy, for large-scale harvesting of information pertinent to launching social engineering attacks. We also show that such harvesting can be undertaken "passively" without resorting to invasive measures, enabling vulnerability assessments which do not rely on exercising deception during social engineering penetration tests. Through automated identification of OSINT that may be used to conduct or enhance a social engineering attack against an organisation, we aim to highlight potential risks to the target, allowing appropriate mitigation techniques to be selected.

The key contributions of our work are as follows:

- In-depth interviews were conducted with expert social engineering penetration testers to better understand the variety of social engineering attacks used, and how OSINT data facilitate the attacks. A summary of the valuable insights from these interviews is presented in Section 3.
- We present an automated approach for identifying the employees of an organisation from amongst the many connected profiles in online social networks. So far as we are aware, no previous work exists on the topic of automatically identifying – from only public data – which of an organisation's social media followers are actually its employees. The nearest approximation we are aware of is Scheelen et al. (2012), who investigated a single company by connecting with followers on LinkedIn, where the social media structure is based around employment.
- We present an approach for automated resolution of identities across social media – demonstrating that large-scale harvesting of such information is feasible for attackers. While employees may be careful about their presentation in online profiles linked to their work identity, we identify features that can be used to link profiles on different online social networks. We present an ensemble classifier, which makes its decision about whether two profiles can be matched based on the reported matches of sub-classifiers working on specific profile features. While more advanced methods exist which could produce more accurate comparison results for each feature, we employ unsupervised methods which release us from the requirement of obtaining training data for the subclassifiers and which are relatively computationally inexpensive.
- We provide an analysis of the online footprints of 13 critical infrastructure companies, demonstrating the extent of their vulnerability to social engineering attacks based on employee information in online social media. We discover that material sufficient to launch sophisticated email and phone attacks targeted at employees is automatically reachable for all but one of the examined organisations.
- We propose a number of mitigation strategies and make our approach – an automatic social engineering vulnerability scanner – available for organisations to counter such risks.[1]

The rest of this paper is structured as follows. In Section 2 we discuss related work connecting OSINT and social engineering. In Section 3 we summarise the findings from in-depth interviews with social engineering professionals. In Section 4.1 we demonstrate how automated methods can be deployed to identify a company's employees from amongst its followers on Twitter, while in Section 4.2 we detail and evaluate our probabilistic identity resolution system on profiles from across four major online social networks (OSNs). In Section 5 we go on to present the results of automated analysis on the digital footprints of critical infrastructure organisations. Section 6 presents the final product of the research as a vulnerability scanner and mitigation tool, evaluating its performance with five companies. In Section 7 we discuss our results and reflect on the implications for social engineering penetration testing and organisational practices for online security. We draw conclusions and offer suggestions for future work in Section 8.

[1] Available from: https://github.com/Betawolf/social-vuln-scanner.