On The Threats To Cloud-based Online Service Users (And What We Can Do About Them)

Gianluca Stringhini
University College London
London, UK
g.stringhini@ucl.ac.uk

EXTENDED ABSTRACT

INTRODUCTION

The Cloud has made it possible for online services of unprecedented size to exist, such as online social networks, webmail-based services, and cloud storage. The largest online services, such as Facebook and Gmail, have hundreds of millions of users. Such users use online services to store their files, stay in contact with friends and colleagues, and much more. Unfortunately, such popularity attracted malicious users too. Cybercriminals use online services to spread spam, perform phishing attacks, steal sensitive information from their victims, and much more.

In this paper, we analyze some of the threats faced by Internet users while using online services. We identify two main threats to users: one posed by large-scale malicious activity (e.g., botnets) that uses the online service to spread malicious content or collect sensitive information, and the other one posed by targeted attacks that compromise user accounts (for example on online social networks) and use their reputation to spread false information.

For each of the two types of threats that we identify, we present some countermeasures that we developed to detect and block such malicious activity on online services. We hope that these techniques will serve as a foundation for online services to develop better defenses and keep their users safe.

BOTNET ACTIVITY ON ONLINE SERVICES

Botnets, networks of compromised computers acting under the control of the same cybercriminal, are commonly used by miscreants to perform malicious activity on online services. Typically, a cybercriminal will use a number of fake accounts on the service and instruct his bots to contact them. These bots will then be used to spread malicious content (spam, malware, phishing scams) or to crawl the service, collecting sensitive information about users that can be then used to perform additional attacks (for example including a user's personal information as part of a social engineering scheme).

From the management point of view, a cybercriminal who wants to misuse an online service will need two elements: a set of infected computers (bots) to use and a set of malicious accounts that the bots can connect to. Over time, different bots in the botnet will access these malicious accounts. By looking at the set of IP addresses that access a group of accounts under the control of the same cybercriminal, we can identify *communities*. Since having a number of accounts accessed by a common large set of IP addresses is infrequent for legitimate accounts, we can leverage this element to detect malicious activity. We developed a system called `EvilCohort` [1] which is based on this insight, and we tested it on multiple datasets coming from a webmail provider and different online social networks. We showed that `EvilCohort` can effectively detect malicious accounts on online services regardless of the type of malicious activity that they perform.

As an alternative countermeasure, we developed `SpamDetector` [2], a system that looks at characteristics that are typical of fake accounts on online social networks and flags such accounts in real time.

ACCOUNT COMPROMISES ON ONLINE SERVICES

Using fake accounts on online social networks has a limited utility for cybercriminals, mostly because such accounts do not have an established reputation and therefore it is less likely for users to believe to the content that they share. For this reason, cybercriminals who want to perform more advanced attacks, such as spreading false news, routinely hijack reputable accounts for this task. This practice is very effective, as it is shown by the compromise of the Associated Press Twitter account in 2013, which announced a terrorist attack against the White House and had a measurable impact on the stock market.

To detect compromised accounts on online social networks, we leverage a simple insight: users develop habits when interacting with social media, and these habits hardly change over time. Conversely, and attacker getting access to an established account and posting malicious content is likely to show differences in his behavior compared to the typical one. Based on this observation, we developed `COMPA` [3], a

system that learns the typical behavior of online social network users and flags as anomalous any message that does not match this behavior. We tested `COMPA` on multiple high profile compromises that happened on Twitter over the last years, with very good results. In addition, we collaborated closely with Twitter and helped them detect more than 300,000 regular user accounts that had been hijacked and were used by cybercriminals as part of spam campaigns.

Online social network accounts are not the only ones that get compromised by cybercriminals. Corporate email accounts are often hijacked too, and used to perform *spearphishing* attacks. To fight this threat, we developed a system, called `IdentityMailer` [4]. Similar to `COMPA`, this system learns the typical behavior of email users, and flags emails that do not match this behavior as possible spearphishing attempts. Because of the longer content contained in emails compared to online social network messages, `IdentityMailer` makes a heavy use of stylometry to perform its detection. We tested the system on real spearphishing emails, with promising results.

CONCLUSIONS

In this paper we highlighted some threats that affect the users of cloud-based online services. We then presented a handful of systems that we developed over the last years to detect and block malicious activity on online services. Although the threat landscape is constantly evolving, we believe that such techniques constitute a good foundation to increasing the security of online service users.

BIOGRAPHY

DR GIANLUCA STRINGHINI is an assistant professor in the Departments of Computer Science and Security and Crime Science at University College London. He obtained a PhD in Computer Science from UC Santa Barbara in 2014. During his PhD, he was awarded multiple prestigious prizes such as the Outstanding Dissertation Award from the Department of Computer Science at UCSB in 2014, the Symantec Research Labs Graduate Fellowship in 2012, and the Best Student Paper Award from ACSAC in 2010. Dr Stringhini has published in top tier computer security conferences such as the USENIX Security Symposium, the ACM Conference on Computer and Communications Security, and the Network and Distributed Systems Security Symposium.

REFERENCES

[1] G. Stringhini, P. Mourlanne, G. Jacob, M. Egele, C. Kruegel, G. Vigna. "EvilCohort: Detecting Communities of Malicious Accounts on Online Services," in USENIX Security Symposium, 2015.

[2] G. Stringhini, C. Kruegel, G. Vigna. "Detecting Spammers on Social Networks," in Annual Computer Security Applications Conference (ACSAC), 2010 .

[3] M. Egele, G. Stringhini, C. Kruegel, G. Vigna. "COMPA: Detecting Compromised Accounts On Social Networks," in Network and Distributed Systems Security Symposium (NDSS), 2013.

[4] G. Stringhini, O. Thonnard. "That Ain't You: Detecting Spearphishing Through Behavioral Modelling." Conference on Detection of Intrusions and Malware and Vulnerability Assessment (DIMVA), 2015.