How long does it take to get owned?

Authors
David Wardle, MSc (Royal Holloway, 2018)
Jorge Blasco Alis, ISG, Royal Holloway

Abstract

This article is based on my thesis - *How long does it take to get owned?* - a study to investigate the amount of time that it takes for stolen credentials to be used by a hacker. It describes the design of fake online identities ("honey identities"), an infrastructure for monitoring their activity and an experiment to test the developed prototypes in which we discovered that it can take as little as 34 minutes for a leaked password to be used.[a]

[a] This article is published online by Computer Weekly as part of the 2019 Royal Holloway information security thesis series `https://www.computerweekly.com/ehandbook/How-long-does-it-take-to-get-owned`. It is based on an MSc dissertation written as part of the MSc in Information Security at the ISG, Royal Holloway, University of London. The full thesis is published on the ISG's website at `https://www.royalholloway.ac.uk/research-and-teaching/departments-and-schools/information-security/research/explore-our-research/isg-technical-reports/`.

Introduction

> *"The reuse of passwords is the number one cause of harm on the internet"*
> Alex Stamos, former CSO of Facebook, Web Summit 2016

A leaked password can have serious consequences if it is re-used across a number of different services, yet password re-use is still a very common practice. This has allowed credential stuffing to become a serious threat. Credential stuffing attacks automate the process of logging into another ("un-breached") website using stolen credentials and is one of the most common techniques used to take control of a user account. An attacker can then syphon the compromised account of its stored value, financial information and other personal information.

The main objective of my MSc project was to design and develop a framework that would allow us to measure the amount of time that it takes for a hacker to log into an account after receiving stolen credentials. The framework also allowed us to monitor attempts to access accounts for different services using the same password or, if the main email has been compromised, reset the password for other services.

Honeypots and honeytokens

A **honeypot** is a security resource whose value lies in being probed, attacked or compromised.

A **honeytoken** is a honeypot which is not a computer, for example a credit card number or special login.

Honey identities

In order to facilitate an experiment to investigate the use of stolen credentials, we developed the concept of "honey identities". Honeypots, and the related honeytokens, have been widely used in network defence since the early 1990s and today there are a huge number of tools available that can complement intrusion detection systems (IDSs) and firewalls as well as detecting malware and spam.

We believe we are the first to come up with the idea of a *honey identity*, which can briefly be described as the collection of personal data used to create a number of accounts (or a digital footprint) whose value lies in being attacked or compromised.

Honey identities

We tried to design the honey identities so that they were both realistic and attractive to any potential hackers. At a high level, it was relatively easy to give the characters a degree of believability by randomly selecting realistic attributes (such as name, gender, etc.). However, it would prove more difficult to maintain the façade should a hacker gain access to one of their accounts or indeed separate accounts for two or more honey identities. "Attractiveness" was also a difficult property to obtain; it was necessary to not only encourage hackers to gain access to a honey identity's account but to then perform further activity within the account and even compromise other accounts belonging to that honey identity - all without revealing the honey identity's true purpose. We speculated that creating a digital footprint consisting of accounts for several high-value websites would make the honey identity "desirable" to attackers. However, this, in turn, is a very subjective quality and in the end we used a combination of reported prices for stolen credentials and website popularity to create a shortlist of candidate services.

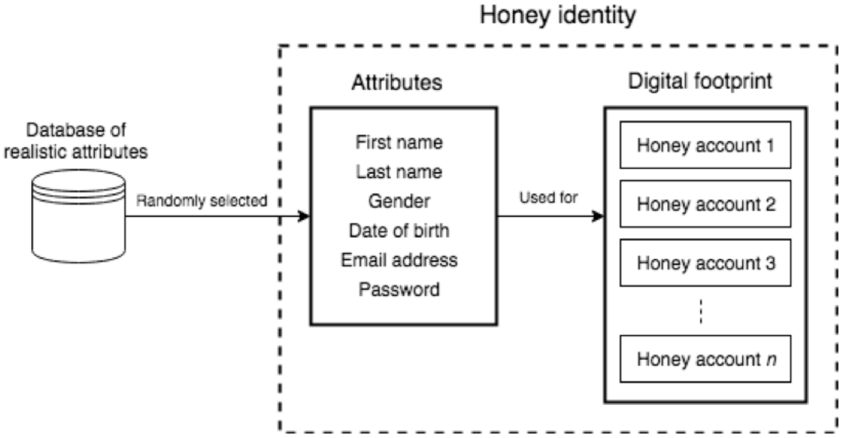

Monitoring infrastructure

Honey identities are only effective if there is some way to monitor their activities. Due to the nature of the accounts, any access/activity to them can be classified as unauthorised, therefore, we also designed a monitoring infrastructure to work alongside our honey identities. We identified several different techniques that allowed us to observe logins and account activity, with varying levels of details:

- Using security functionality provided by the website, such as login notification emails and Two Factor Authentication (2FA) SMS messages.
- Observing account activity (e.g. outgoing/incoming email, searches, etc.).
- Checking the credentials were still valid.
- Hiding honeytokens within the account.

Some of these methods were a lot more powerful than others. For example, a 2FA text message would alert us that someone had attempted to log into an account but would rarely provide any more information such as IP address or even the username of the account. However, by using a combination of these techniques, such as login notification email along with account activity and honeytokens, it was possible to get a lot of information about any hacker and, in some cases, get an insight into their motivation. Were they looking to hijack the account, use it for spamming purposes, searching for sensitive information, or just curious?

One avenue explored, but ultimately proved unfeasible for the MSc project, was a partnership with web services. This could have allowed us to obtain extensive logs for access as well as helping with bulk account creation. To simulate this, we created a website and email server for a fictional financial services company.

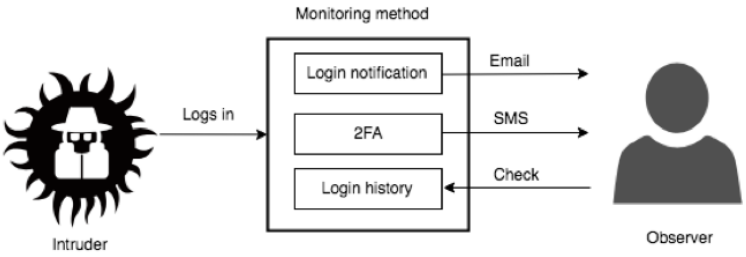

Prototype

After designing the honey identity and monitoring infrastructure, we created simple prototype versions for both. The size of the honey identities was limited to just six public websites along with email and website accounts for our fake company. We developed software that would generate various attributes/meta data for the identity, use this information to automatically register accounts on our chosen websites and store the resulting credentials in a local database.

On the monitoring side of things, all of the honey accounts were configured so that any security functionality available would provide as much information as possible about any activity within the account. However any login protection features were disabled since we wanted any hacker to gain full access to the accounts. For simplicity, all email accounts were set up to forward to a separate, single address. A webservice was also used to forward SMS messages to the same email address.

We populated every email account that we created with a collection of fake emails, some of which contained honeytokens or further credentials for other accounts belonging to that honey identity. This added a degree of realism to them as well as helping with monitoring.

How long does it take to get owned?

We released the credentials for a number of honey identity accounts on a popular paste website - emulating a common method that cyber criminals use to share stolen credentials. The format of the leaks varied for each honey identity, some passwords were displayed as plaintext whilst others were the result of a strong password hash algorithm. However, the hashed passwords were relatively weak and would have only taken a few seconds to crack with a common wordlist.

Paste website

A service that allows users to store and share plain text (such as snippets of code). Often users can post anonymously which has given rise to the sharing of credential dump.

Unsurprisingly, all of the accounts that had their passwords exposed in plaintext were accessed at some point during a month-long observation period. The fastest access occurred just 34 minutes after the login details were published. Even though it would have been trivial to crack the password hashes, none of these credentials were used.

Even with a limited number of results, it was possible to see some interesting trends. In all but one case, the hacker's IP address was traced to eastern USA. It is highly unlikely that this was their real location and it is very easy to fake this information through the use of a VPN or proxy server. Indeed, a search on SPAMHAUS (a database of known spammers) revealed that one of the IP addresses was part of a malicious botnet. Masquerading as an American web user has two benefits for a hacker: it hides their real location and lowers the likelihood of their login being noticed by pro-active service providers.

After looking at the various activity, we speculated that all of the access was conducted by real people rather than software to validate exposed credentials. Throughout this article, I have used the term “hacker” for those that used any of the leaked passwords, however, I believe that every login was by someone who was simply curious rather than having any malicious or criminal intent. We could see that one user browsed the fake company website and read the victim honey identity’s profile. Another person deleted a honey identity’s entire google account shortly after logging into it. Thankfully, it was relatively straightforward to recover the account and we discovered their search history included “how to delete gmail account” giving us an insight into their level of technical expertise.

To be honest, the type of people accessing our honey accounts was not a surprise. In fact, the experiment relied on curious people to find the credentials themselves, either by browsing recently posted posts on the paste website or using a search engine dork (a very specific search term). It is also likely that any “l33t” (elite) hackers or cybercriminals would have had strong suspicions about the true nature of our published credentials. It is important to remember that the experiment was not the main objective of the project, rather a means to test the prototype framework/infrastructure that we he had created. In this aspect, we were very pleased with the results that we achieved.

Further work

As mentioned, there is scope to conduct a fuller experiment by advertising the credentials on dark market websites (under the guise of a bigger breach). This would hopefully allow for greater analysis of hacker activity within compromised accounts. A more complete experiment would require hundreds (thousands?) of honey identities, registering on a significant number of websites and simply waiting to be hacked. An experiment of this nature would accurately answer the question posed in the title of my MSc project, *How long does it take to get owned?*. There is also the opportunity to investigate the make-up of an attractive online identity which, in turn, could be used to improve honey identities.

In our experiment, we didn’t witness any case of someone attempting to re-use the credentials on another service, at least not one within the honey identity’s scope. With a longer timescale to conduct an experiment, it may be possible that our leaked passwords would be added to a real set of stolen credentials and provide more research opportunity into the threat of credential stuffing attacks and potential damaged caused by password re-use.

The honey identities that we created for our project were basic prototypes and there is a lot of room for improvement. One area that would significantly improve the realism of the identities would be to give them an active digital footprint. On social media networks, this could be a matter of publishing content on a regular basis, however this brings its own risks of being flagged as spam. Another aspect is the identity’s metadata. So far all of the attributes were completely fake but what if we used a real postal address - would we start to receive nuisance mail? Pre-paid and virtual debit cards would also provide the opportunity to use valid financial details with relatively low risk.

Conclusion

Thanks to regulatory powers, companies have been forced to take information security a lot more seriously over the last few years. However it is often said that there is no such thing as 100% security and large data breaches will continue to occur. If you are an end user, it is crucial that you follow best practice - use a password manager to ensure that every password is strong and unique, and enable multi factor authentication whenever possible. From a technological point of view, systems should be designed and developed to encourage (but not necessarily enforce) users to follow these good practices.

Biographies

David Wardle completed his MSc in Information Security at Royal Holloway, University of London, graduating in 2018 with distinction. David is currently working as a freelance web developer, with over fifteen years experience.

Jorge Blasco obtained his PhD from University Carlos III of Madrid in 2012. His dissertation was focused in the field of information security. After obtaining his PhD, Jorge worked as an assistant lecturer in University Carlos III of Madrid. In 2014, he moved to City, University of London, where he worked until 2016 as a Research Fellow in a project about application collusion. His main research interests include mobile malware, steganography and wearable devices. Since September 2016, Jorge Blasco is a Lecturer and MSc in Information Security Course Director in the Information Security Group at Royal Holloway, University of London.

Series editor: S.- L. Ng