Social Media:

The Overlooked Weak Link in Security

Lexi Walker

December 13, 2019

Abstract

In discussion about security, the “weakest link” is often left out: the user. Users who aren’t security-conscious often overlook things that can leak unintended information to malicious individuals. As more information about each of us is available on the internet, it is becoming even easier to exploit people in this way.

One emerging source of information for social engineering is social networking sites, such as Facebook and Instagram. In this paper, I will describe the kinds of ways in which a person’s social media presence opens them up to information gathering, and therefore, targeted, convincing phishing attacks.

I will also discuss my own experience of crafting phishing attacks for three of my classmates through information gathered from their instagram feeds, and the results of the attacks, as well as a holistic view of the kind of information I was able to gather.

Introduction

The user is an often-overlooked part of the security workflow. When the user is uninformed about good security practices and the potential danger of leaking seemingly-harmless information, it exposes them and their company or organization to social engineering, the practice of manipulating individuals in order to get access to personal or confidential information. Even in its simplest form, social engineering is incredibly powerful; purely through socially engineering a Verizon employee, an attacker was able to get access to former CIA director John Brennan’s personal email, which contained confidential files regarding national security matters.[1]

When we post details about our personal and professional lives online, we greatly expose ourselves to social engineering. This information can be used to get access to your accounts through security questions, or to gain your trust and collect even more information from you, such as a phishing scheme to get online credentials, your credit card number, or even your social security number.

To The Community

This topic is important for visibility among users of social platforms, particularly younger users. As the barrier of getting online continues to decrease, new users are getting online at increasingly younger ages, and with less of an understanding of the possible impact that it could have on their lives. Many of these young people haven’t known a world without the internet, and without being constantly surrounded by technology, which creates automatic trust and acceptance of it. The more we are online, the more we must be careful with whom we’re interacting, and what information we’re sharing; with such a young audience on social media today, we must be sure to educate this audience with the potential implications of their actions.

---

[1] Zetter, Kim. “Teen Who Hacked CIA Director's Email Tells How He Did It.” *Wired*, Conde Nast, 30 June 2017, www.wired.com/2015/10/hacker-who-broke-into-cia-director-john-brennan-email-tells-how-he-did-it/.

In order to protect ourselves, we must be skeptical and alert while online - this means being wary of interacting with individuals that we haven't met offline, and being more conscious of our media presence, and what information about us exists to the general public. While some information about you online can't be controlled, such as articles written, we can curate our own online presence, and consider the impact that an attacker could have given that content.

Background

Human error and lack of judgment are crucial to many data breaches, both small and with global impact. In April 2013, the Associated Press Twitter was hacked, and a tweet reading "Breaking: Two Explosions in the White House and Barack Obama is injured" was posted to 2 million followers, causing The Dow Jones to drop more than 140 points.[2] This event, which impacted a wide range of people, as well as the American economy, was caused by a fake email, in which the attacker, the Syrian Electronic Army, linked the user to a website that prompted the recipient for login details for the AP Twitter account: "That the name in the 'From' field of the email didn't match the name in the signature line was the only clue that the email was fake."[3]

Even RSA, an American computer and network security company, is not immune to social engineering. An attacker sent phishing emails with the subject line "2011 Recruitment Plan" to a small set of employees with an attached Excel file, which contained malware that used a "zero-day...flaw in Adobe's Flash software to install a backdoor."[4] While it remains unclear exactly how much information the attacker was able to uncover, it's likely that other clients of RSA, L-3 Communications and Lockheed Martin, were attacked as a result of the RSA breach.[5]

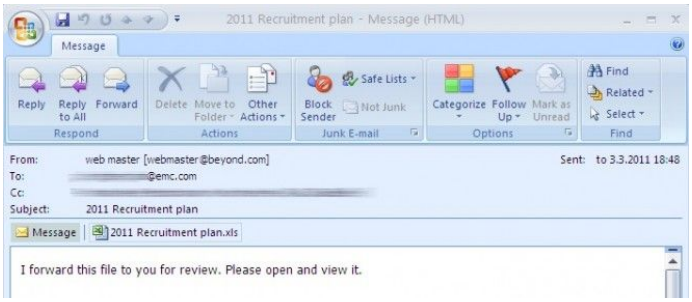

Screenshot of Phishing Email From the RSA Attack[6]

---

[2] The shocking tweet came from the Associated Press earlier this afternoon: "Two Explosions in the White House, and Barack Obama is injured.". "AP Hack Proves Twitter Has a Serious Cybersecurity Problem." *CNNMoney*, Cable News Network, money.cnn.com/2013/04/23/technology/security/ap-twitter-hacked/index.html.

[3] Dobran, Bojana. "Be Prepared: 7 Most Famous Social Engineering Attacks In History." *PhoenixNAP Global IT Services*, PhoenixNAP Global IT Services, 3 Oct. 2019, phoenixnap.com/blog/famous-social-engineering-attacks.

[4] Richmond, Riva. "The RSA Hack: How They Did It." *The New York Times*, The New York Times, 2 Apr. 2011, bits.blogs.nytimes.com/2011/04/02/the-rsa-hack-how-they-did-it/.

[5] Dobran, Bojana. "Be Prepared: 7 Most Famous Social Engineering Attacks In History." *PhoenixNAP Global IT Services*, PhoenixNAP Global IT Services, 3 Oct. 2019, phoenixnap.com/blog/famous-social-engineering-attacks.

[6] Zetter, Kim. "Researchers Uncover RSA Phishing Attack, Hiding in Plain Sight." *Wired*, Conde Nast, 3 June 2017, www.wired.com/2011/08/how-rsa-got-hacked/.

Even without much insider information, these attackers were able to achieve major impact both within the targeted company and beyond due to employee naiveté, as well as sending the attack to a wider audience (multiple employees). With a more specific target, the success rate is more crucial; therefore, the attack must be more convincing.

This concept is called whaling, in which an entire phishing attack is crafted to target a single high-profile target.[7] For an experienced social engineer, a few key details can lead to a huge amount of information gathering; Rachel Tobac, "a celebrity among the DEF CON crowd", was able to get access to a CNN reporter's home address, phone number and hotel points solely through an instagram check-in at a hotel and a tweet about furniture.[8]

With two small social media posts able to access that kind of personal information, I wondered what kind of impact an entire social media feed could have, and what kind of phishing attacks could be crafted from that information.

The Ploy

I decided to target three classmates of mine - one who is currently taking COMP 116 (#1), another who has previously taken COMP 116 (#2), and a third who does not study computer science (#3) - with phishing attacks based on information I could gather from their Instagram feeds. I created a report of the information I was able to gather from posts over the past two years, and sent a targeted phishing email to each of them, which directed them to click on a link provided in the body of the email. At this link, I created a website that announced they had been phished, and directed them to fill out a survey.

Information Gathering

In social media culture, particularly for teens and young adults, there is pressure to have a large number of followers. According to Alexandra Fabugais-Inaba from Rutgers University, "The societal norm has become that more people should care about your life compared to how much you care about others' lives … Defined as the follower to following ratio, your ratio is higher the more followers you have and less people you follow."[9] This creates pressure to accept followers that are either people you don't know, or are content accounts, posting, for example, dog pictures or inspirational quotes; these accounts could be run by anyone. When a user allows these accounts to follow them, these individuals get access to their entire social media feed, regardless of whether their accounts are private or public.

While I did analyze private accounts from my own personal account, I feel confident that I could've created a fake instagram and successfully followed my classmates; each followed either content accounts, or accounts that looked fake. One such account, followed by #2, had

---

[7] "Whaling Phishing Attacks Explained: What Is Whaling?" *Rapid7*, www.rapid7.com/fundamentals/whaling-phishing-attacks/.

[8] O'Sullivan, Story by Donie, et al. "We Asked a Hacker to Try and Steal a CNN Tech Reporter's Data. Here's What Happened." *CNN*, Cable News Network, 18 Oct. 2019, www.cnn.com/2019/10/18/tech/reporter-hack/index.html.

[9] Fabugais-Inaba, Alexandra. "Why You Need to Stop Caring About Your Instagram Ratio." *Study Breaks*, 21 Aug. 2018, studybreaks.com/tvfilm/instagram-ratio/.

the username “run_the_world”, the name “Run the World”, and a single post of a landscape photo with no description. Additionally, #2 didn’t follow this account back. Altogether, this is very likely a fake or bot account, which has access to all of #2’s information. Additionally, #1 has an account “ant_mitchelll” following her, who she admits she doesn’t know, and who is following many more people than are following them - this seems suspicious and possibly malicious.

I was able to gather the most intel on #1. She posted a lot of concrete information, such as tagging locations, elaborate captions, and images of specific items or events. From looking through photos, captions, comments, and location tags, I was able to gather her sorority, university, hometown, major, club and position within the club, sibling’s name, school, and graduation year, internship company and location, and birthday.

Next was #2; while not quite as open, I was able to find her university, vacation location and rough dates, major, summer location, sorority, an internship from two summers ago and the title of the research she did there, and hometown.

Lastly, while #3 was very active on social media, I wasn’t able to gather as much concrete information; I was able to find his university, home state, graduation year, boyfriend’s name and birthday, study abroad location and approximate dates, interest in theater, and previous employment. Given the number of posts, I was expecting to find a lot more than I did.

The Phishing Emails

Based on the information I gathered, I created possible phishing scenarios for each target, and selected the one I thought would be most convincing.

For #1, I saw that she was passionate about women in computer science; she had worked for Girls Who Code, spoke at a women empowerment conference, and was an organizer of a Women in Tech conference. I decided that this angle, which would be emotional, might distract her from an email that doesn’t quite add up. I decided to play the angle of a student at another university trying to create their own Women in Tech conference, and reaching out for help. I decided on NYU, a school from her home city, and looked up their Women in Tech club. They have a group called Women in Computing (WinC), and the organization has an executive board. I decided to impersonate Ishita Mittal, their current Event Coordinator, and created a Gmail with her name - I used the username ishmittal1999 - she’s a Junior, so I thought adding 1999 (likely her birth year) might help to sell it. I also created a GitHub with the username “nyuwit”, on which I hosted an alleged website for the new NYU conference.

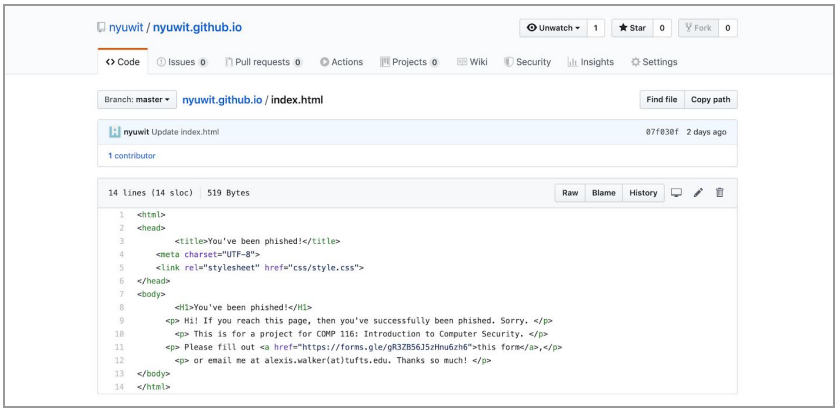

I then sent an email to #1 from the fake Ishita email, claiming to be trying to start up a new Women in Tech conference at NYU, asking her to take a look at the preliminary conference website, which was actually linked to the phishing site.

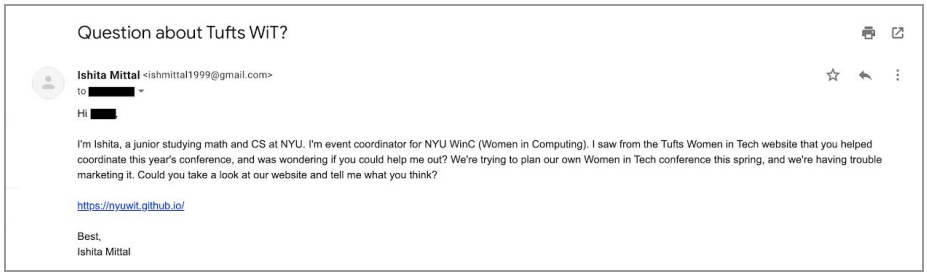

For #2, I saw that she had done research in the community health space for a summer internship (she posted a picture with a research poster), and that she was wearing a Peer Health Exchange shirt; Peer Health Exchange is a group focused on “empower[ing] young people with the knowledge, skills, and resources to make healthy decisions”.[10] Seeing that she was involved with this kind of work for both an internship and as an extracurricular in her free time, I thought this might similarly hit something she’s very passionate about, and therefore allow her to overlook anything that might seem weird about the email or the link.

I decided to impersonate someone working under a professor working on similar work. I confirmed that this information about her research was available publicly; sure enough, the title of her poster and a brief explanation of her work was there. When I searched the title of her poster, I found a paper on a similar topic co-authored by a professor at the University of Leiden, Nadia Garnefski. I found a link to a piece of work she is currently working on, still related to #2’s research. I knew that a GitHub.io link would never be convincing for asking about something related to community health research, so I decided to use the same github as for #1, but hyperlink it over a real link to this research. I also decided to use the same email, as I thought pretending to be an undergraduate research assistant was much more realistic than if the professor herself reached out. The final email asked #2 if she would be willing to chat about her

[10] “About Us.” *Peer Health Exchange*, www.peerhealthexchange.org/about-us.

work at the NIH (where she had interned), and to take a look at the link for reference on the work that was ongoing at Leiden.

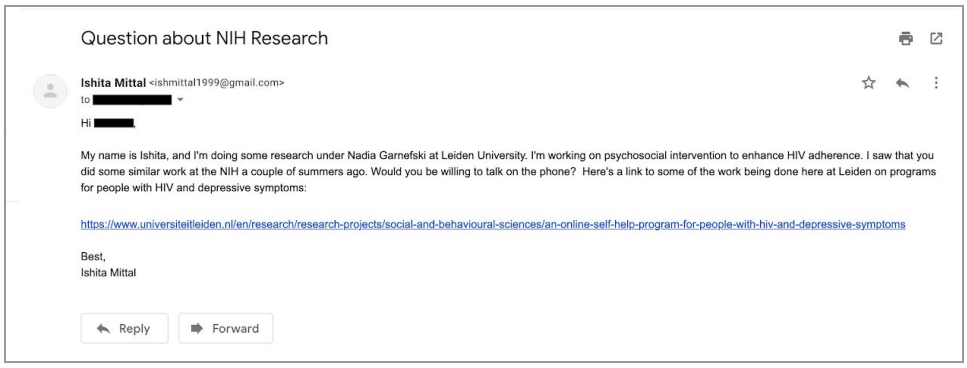

Question about NIH Research

Ishita Mittal <ishmittal1999@gmail.com>
to [redacted]

Hi [redacted],

My name is Ishita, and I'm doing some research under Nadia Garnefski at Leiden University. I'm working on psychosocial intervention to enhance HIV adherence. I saw that you did some similar work at the NIH a couple of summers ago. Would you be willing to talk on the phone? Here's a link to some of the work being done here at Leiden on programs for people with HIV and depressive symptoms:

https://www.universiteitleiden.nl/en/research/research-projects/social-and-behavioural-sciences/an-online-self-help-program-for-people-with-hiv-and-depressive-symptoms

Best,
Ishita Mittal

Reply Forward

Lastly, I had a bit of difficulty with #3. I wasn't confident that I could make a very targeted phishing attack based on the information that I gathered. One information did catch my eye, though: his bio says "Tufts '20", meaning he's either graduating in the spring or next fall, though much more likely the spring. This means that there are forms due all the way from October until May; I decided to send a false "urgent reminder" about a form due soon.

Again, I knew the form link had to look like it was coming from a university URL, so I used the same GitHub link tagged over the actual Tufts graduation checklist. I also didn't think that I could pass off as an administrator coming from a Gmail account, but might be able to pull off being an undergraduate administrative assistant. I BCC'd #3, to make it look like multiple people might have received it.

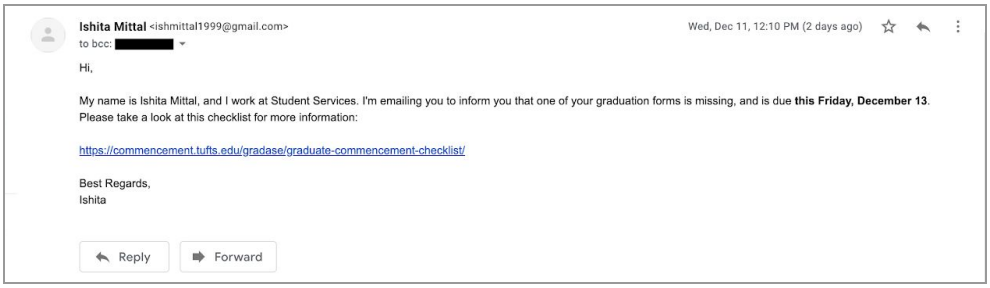

Ishita Mittal <ishmittal1999@gmail.com> Wed, Dec 11, 12:10 PM (2 days ago)
to bcc: [redacted]

Hi,

My name is Ishita Mittal, and I work at Student Services. I'm emailing you to inform you that one of your graduation forms is missing, and is due **this Friday, December 13**. Please take a look at this checklist for more information:

https://commencement.tufts.edu/gradase/graduate-commencement-checklist/

Best Regards,
Ishita

Reply Forward

All three email links led to this site:

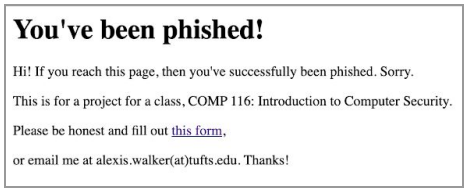

You've been phished!

Hi! If you reach this page, then you've successfully been phished. Sorry.

This is for a project for a class, COMP 116: Introduction to Computer Security.

Please be honest and fill out this form,

or email me at alexis.walker(at)tufts.edu. Thanks!

Results

#1 was phished successfully. In her feedback, she writes, "Since it was during reading period, I wasn't sleeping very much, so I wasn't very alert. I think if it had been a normal week, I would've looked at the email more carefully."

#3 was also successfully phished. In his feedback, he writes, “with the chaos of senior year, this was a very smart way to phish someone.”

#2 was not successfully phished, but was close. #3 had asked #2 whether she had received a weird email before she got the chance to see my email; when she checked her email, she was on higher alert for anything possibly suspicious. If she hadn’t received the warning, she said, she likely would’ve clicked on it. The network effect is crucial; within an organization or social group, one person who received a scam might warn others before they open it, though they wouldn’t have thought it was suspicious otherwise. I think sending the attacks so close together was a mistake; if I had the opportunity to do it over again, I would’ve spaced them apart. In this case, #2 may have forgotten about #3’s warning by the time she received the email, or may have received hers first.

Conclusion

Overall, I was surprised that these whaling attacks were so successful. I’m very bad at deception, and yet, two out of three attacks were successful. While proud, I’m also a little concerned at the prospect of what an experienced social engineer could do with the same information, or even someone of my level with malicious intent. When something is posted on the internet, it’s nearly impossible to fully erase it; Before young people start using the internet, and especially before they start using social media, it’s important that we communicate these potential risks.