Article development led by **acmqueue**
queue.acm.org

Investigating the emerging black market of retail email account hacking services.

BY ARIANA MIRIAN

Hack for Hire

A SINGLE EMAIL address often underpins one's entire online identity, from banks, to business, to social media profiles and more. This identity is used not only when registering for this multitude of services, but also when passwords for these services must be reset. Thus, an attacker gaining access to an email account poses the risk of compromising all the other services tied to that account as well. Politicians, journalists, and cryptocurrency folks have all been the victims of targeted attacks that started with access to their email accounts that then wreaked havoc on other online accounts tied to those email accounts.^{3,7,9}

Since email accounts can provide a wealth of information, many types of attacks target them: password guessing, access token theft, password reset fraud, and phishing, to name a few. Email providers have added mechanisms such as security questions, spam filtering, and two-factor authentication to limit the success rate

of these attacks.^{4-6,12} These defenses prevent many compromises, thus increasing the sophistication and time needed to access accounts. Targeted attackers, however, are willing to put in the extra effort needed to access an account in the face of these large-scale defenses.

While targeted attacks are often thought of as requiring nation-state capabilities, there is an emerging black market for “hack-for-hire” services, which provide targeted attacks to anyone willing to pay a modest fee. These services purport to be able to break into the accounts of a variety of different email providers, an example of which is shown in Figure 1. As these services are just emerging, little is known about how they attack their victims and how much of a risk they pose.

To understand this risk, we investigated the hack-for-hire black market, identifying 27 retail email account hacking services and purchasing these services from them. Using covert identities, we engaged with these services to break into purported “victims”—in truth, Google accounts that we controlled. Working with Google, we recorded both our interactions with the hijackers and how these hijackers tried to attack our victims.

To Catch a Hijacker

As a whole, the targeted hijacking black market was riddled with scams, but a handful of services launched sophisticated attacks that leveraged phishing as their main attack vector. These attacks were persistent, personalized, and able to bypass SMS two-factor authentication (2FA). Using signals derived from these attacks to identify other victims, we estimate that attackers target about one in every million Google accounts. Given the sophistication of the phishing attacks, we believe that the best line of defense for at-risk users is to protect accounts with universal 2nd factor (U2F) security keys as a 2FA mechanism. U2F security keys protect against sophisticated phishing attacks because the U2F protocol validates the domain before

sending the 2FA code, preventing a user from getting phished. Though the focus of this investigation was on Google accounts, the lessons learned generalize well across email providers.

Discovery of services. The investigation of hack-for-hire services began by searching English, Chinese, and Russian black market forums for advertisements related to targeted account hijacking. We also searched Google for hijacking-specific keywords to identify services with public-facing storefronts, and we contacted the abuse teams of large Internet companies for leads on any such services they were tracking. In all, 27 prospective hack-for-hire services were identified. The majority of these services advertised in Russian, and ranged anywhere in price from \$23 to \$500 per contract.

Posing as a buyer. For every hack-for-hire service contacted, we communicated via a unique “buyer persona” to pro-

tect our identity and to avoid linking our interactions across services. Each persona involved selecting a name in the native language of the hack-for-hire service. For example, if the service advertised in Russian, then we chose a common first and last name in Russian for our persona. We also created a Google account for each persona to use for all email communication. For non-English services, a native speaker performed all translation when communicating.

Selecting a victim. When contracting hack-for-hire services, we created a victim persona to serve as a target. The victim persona was given a large digital footprint to craft a realistic online presence. This meant creating a name and Google account for the victim in a similar fashion to the buyer persona. The victim persona’s inbox was populated with a subset of messages from the Enron email corpus to give the impression that the Google account was in ac-

tive use.² We replaced the names and other identifying information from the Enron messages with the victim’s information. Moreover, the victim persona’s Gmail address was protected by SMS 2FA, the most widely used form of 2FA today.¹ This was used to determine if the hack-for-hire services would be able to bypass this type of protection.

In addition, we created a Web page that advertised a small business that the victim either owned or worked at. We purchased the domains of the Web pages from auction to ensure each domain had prior history. We also purchased privacy protection for each of the domains to protect the registration information (one recent study showed that 20 percent of domains are protected in this fashion, so we did not expect privacy protection to raise any red flags⁸). This webpage linked to the victim’s email address, as well as a fictitious associate’s email address. In this way we could determine if

the hack-for-hire services would attack the associate as a way to gain access to the victim. We also created a Facebook page for the victim to see if the hack-for-hire services would use it in their attacks. All items on the Facebook page were private (a public user would not be able to see these items) except the About Me section, where the victim's Web page was listed (as a personal advertisement for the victim's business).

Monitoring attacks. Because of uncertainty over which attack methods these hack-for-hire services would use, we created an extensive monitoring infrastructure that would inform us when an unauthorized user entered and modified the Google account. We also monitored the websites to record all visitors. This monitoring infrastructure contained: a Google app script for Gmail in every account; Google logs; and a network capture of all traffic to the fictitious websites.

The Google app script loaded into Gmail for each account was a modifica-

tion from one used in a previous study.¹¹ The Google app script would send information to a server controlled via a proxy indicating whether the script was still connected to the account and whether there were any changes to the account. For example, the Google app script would indicate whether a new message had appeared in the inbox or spam folders, if a message was moved to trash, or if any messages marked as unread were read by a user. This logging recorded the actions of the attackers once they were in the account.

We also were able to analyze any login activity to our victim personas' Google accounts. These logs, captured and analyzed by our Google colleagues, recorded login attempts into the account and their origins, brute-force attempts, and whether 2FA was triggered on the account for a suspicious login attempt.

Finally, we captured all network traffic to each victim persona's website.

As mentioned previously, the website address was linked on the About Me section of the victim's Facebook page. If an attacker was able to find the website via the Facebook page, this would be apparent in the network capture. The network capture would also show if any of the hacking attempts to the Google accounts were from the same IP addresses that were visiting the Web pages.

Legal and ethical considerations. Since this study involved engaging with actors that were performing illegal activities, there were legal and ethical questions to consider.

Legally, there were two concerns: unauthorized access into Google accounts and violating Google's terms of service. In the U.S., as in many other countries, unauthorized access into an electronic account is illegal. Hiring services to perform this act could be considered aiding and abetting; however, since the email accounts were directly under our control and we were acting in collaboration with the service provider (Google), we were explicitly authorizing entities to access our accounts. Moreover, creating fake Google accounts violates Google's terms of service, but this study was approved by both Google and the general counsel for UC San Diego, before the work was started.

Although this study is not considered human rights research by our institutional review board because we were measuring organizational behaviors and not behaviors of an individual, there were other ethical considerations. By creating fictitious victim and buyer personas, we removed the possibility of any individual being harmed during this study. Moreover, we interacted with these services within the scope of their terms and paid them if they were successful. We believe the lessons learned from this study outweigh the cost of supporting these services by paying them.

Hack-For-Hire Playbook

The controlled experiment and logging infrastructure allowed examination of the playbook that attackers use to take over a victim account. Only five of the 27 services we contacted actually attempted to break into the victim's Google account. Note that the "success" of a hack-for-hire service was dependent on our actions: in some cases,

Figure 1. Example of a hack-for-hire advertisement.

Figure 2. Lures used to try to access a victim account.

we provided the password or 2FA code when prompted by a phishing page, while in other cases we did not in order to see how the service would adapt.

Overall, we never observed any brute-force login attempts, communication with the victim's Facebook account, or communication with the associate's email. Of the five services that attempted to gain access to the account, one sent the victim malware executable via an email message. While we were not able to run the executable in a virtual machine sandbox, VirusTotal reported the malware executable was a remote access trojan. The other four services used phishing as their primary attack vector. I share our key findings here, but more details can be found in the full article.¹⁰

Email lures and phishing. All of the attacks started with an email lure to the victim account. These lures impersonated some trusted or authority figure, presumably to spur the victim into clicking on the link. There were five types of lures across all phishing messages: those impersonating an associate persona, a bank, a stranger, a government entity, or Google (as illustrated in Figure 2). The associate lures leveraged trust to get the user to click on an “image” (which led to a phishing page), while the stranger lures consisted of an unknown person emailing the user with an “image” or link. However, the government, bank, and Google lures conveyed a sense of urgency in their messages. Figure 3 shows examples of a government lure (translated into English) and a Google lure.

On average, attackers sent 10 messages over the course of 25 days and used different lures during their persistent attacks. Figure 2 illustrates this behavior, showing the time since the attackers sent their first email message, the type of lure for each message, and when we clicked on the lure (potentially halting any future attempts). An X indicates when we clicked on a link in a message sent to a victim. Each row corresponds to one victim and numbers on the right denote the total number of emails sent by a service. The most popular lure mimicked Google, followed by associates, then lures from strangers.

Of the services that sent personalized lures, all but one asked the buyer persona for more details ahead of time (such as the name and email address of

a known associate). One service was able to construct personalized lures without additional information from the buyer persona, indicating that this service searched and found the online website constructed for the victim. We also purchased two other contracts from this service, however, and in both cases the operator asked the buyer persona for details on the victim. These differences in behavior suggest these services have multiple people working behind them.

As mentioned previously, all but one of the services relied on phishing as their main attack vector. Clicking on the phishing link took us to a landing page that looked like the Google login page. After entering our password, we were taken to a page that prompted us for our 2FA code. All of the services that were able to access the account did so by phishing the SMS 2FA code, bypassing this security protection.

Live adaption. While most of the services accounted for the 2FA code in their initial phishing flow, two exhibited phishing attacks that adapted to

obstacles. These two services, in their initial phishing flow, did not account for a 2FA code. Their phishing flow collected the password and then tried to log into the Google account, which was successfully blocked by 2FA. After realizing this, both services sent the victim additional email messages with a different structure from the ones previously sent. One of these two services sent a new message that, when clicked, would request both the password and the 2FA code. The other service also changed its flow to account for the 2FA code, and in doing so sent phishing messages that looked similar to those sent from a third service. These similarities suggest the use of common tools among services.

Note that these services were able to bypass 2FA because they phished the SMS 2FA code from the user. As noted earlier, SMS 2FA is the most widely used form of 2FA, and we wanted to see if those who use it would be susceptible to a hypothetical hack-for-hire attack. To prevent hack-for-hire services from gaining access to an ac-

Figure 3. Examples of a government lure and a Google lure.

Figure 4. Purported prices to access various accounts.

Target	Service A	Service B	Service C	Service D*	Service E*
Mail.ru	\$77	\$77	\$62	\$54	\$77
Rambler	\$152	\$108	\$77	\$77	\$108
Yandex	\$106	\$108	\$77	\$77	\$108
Gmail	\$384	\$385	\$92	\$77	Negotiable
Yahoo	\$384	\$231	\$92	—	—
Facebook	\$306	—	—	—	—
Instagram	\$306	—	—	—	\$231

Figure 5. Monthly prices for service A.

Figure 6. Accounts associated with hack-for-hire services.

count, at-risk populations should use a security key as their 2FA protection, as that code is unphishable.

Post compromise. Upon gaining access to a victim's account, hack-for-hire services start to remove any evidence of compromise and ensure their ability to regain access if needed. Services that gained access to our victim accounts proceeded to sign in to each account and remove all Google email notifications related to a new device sign-in from both the inbox and trash. None of the services changed the password, but we did observe three services remove the 2FA authentication and recovery number from our victim accounts quickly after they gained entry. We presume they took this step to ensure the buyer could gain access to the account and so that the service itself could regain access to the account, but we did not see any service log into the account after the initial login. In essence, the services took precautions to remove their digital footprint from the Google accounts they were breaking into.

Once the email account was accessed, all but one of the services initiated a portability feature in Google called Takeout, allowing them to download the victim account's email content, and provided the parcel of information to the buyer persona. Only one service avoided logging into our victim account and provided the password to the buyer persona without using it first. These findings highlight an emerging risk with data portability and regulations around streamlining access to user data. While intended to improve usability for users, capabilities such as Takeout also increase the ease with which a single hijacking can expose all user data to a service. Since this study, Google has added more step-up verification on sensitive account actions.

Real Victims and Market Activity

Based on our findings from this process, we analyzed the forums of the most successful services to understand their pricing for other services. Moreover, we present an estimate of the number of real victims affected by these services based on login traces from Google. Our findings suggest this market is quite niche.

Alternative services and pricing. While our investigation focused main-

ly on Google—because of legal constraints—many of the hack-for-hire services we interacted with also purported to be capable of breaking into other types of accounts. Figure 4 shows the prices of the hack-for-hire services as of Oct. 10, 2018. All prices are in U.S. dollars, converted from rubles. An * indicates that the service's advertised price was lower than the final payout requested.

Across these five services, hijacking Russian email providers was the least expensive offering, while hijacking a Google or Yahoo account was the most expensive. Breaking into a social media account fell in the middle of these two extremes. The advertisements for one of the services exhibited prices that changed over time, shown in Figure 5. The price of Google account hacking increased the most—from \$123 to \$384 per account over two years—while the cost of Russian email provider hacking has remained largely the same. These differences and price changes are probably the result of a multitude of factors such as demand, changes in security, and competition from other services.

Victims over time. Of the 27 unique services we contracted, only three were able to log into the victim email accounts successfully. Google analyzed associated metadata with the successful login attempts and found that all three services relied on an identical automation process for password validity checks, bypassing any security obstacles such as 2FA, and downloading the victim email archive. While the email sender and delivery addresses differed among the various contracts, the login automation process was the same across the eight months these services were contracted. Google was able to create a signature for this automated login fingerprint and retroactively analyze how many Gmail accounts had a suspicious login attempt.

Google identified 372 accounts targeted by this automated login framework from Mar. 16, 2018, to Oct. 15, 2018, or about one in every one million Google users. Figure 6 shows the weekly breakdown of the number of targeted Google accounts. Be aware that these numbers are lower bounds, since we cannot infer how many users were targeted by these services but did not click on the link (or provide their infor-

mation to grant access), only how many users had an account that was accessed by these services. Despite these limitations, the volume of activity for hack-for-hire services is quite small when compared with other services such as off-the-shelf phishing kits, which impact more than 12 million users a year.¹³ We suspect the hack-for-hire market is small compared with other markets, such as malware distribution.

Discussion

Overall, hack-for-hire services charging \$100–\$400 per contract were found to produce sophisticated, persistent, and personalized attacks able to bypass 2FA via phishing. The demand for these services, however, appears to be limited to a niche market, as evidenced by the small number of discoverable services, an even smaller number of successful services, and the fact that these attackers target only about one in a million Google users. Moreover, this market suffers from poor customer service, as many of the services were slow or inconsistent in their responses to our buyer personas.

Regardless of the behavior of the market, this study sheds light on the importance of security keys for populations who believe they are at risk, as only a security key can protect a user from the attacks viewed in this study. As the market evolves and defenses change, however, attacks might also change and shift from phishing to more persistent threats such as malware.

In conjunction with this study, Google introduced two new defenses to help protect against man-in-the-middle phishing, which in turn would protect against these services. Google now runs additional heuristics when you log in, and also prevents some forms of automated login frameworks. In addition, two of the services have nearly doubled the price of hacking Google accounts since Google rolled out the new protections to users, although it is not known if this price hike is coincidental or was caused by the increased Google protections.

Acknowledgments

Thanks to the co-authors of the original research publication for their feedback in writing this article: Kurt Thomas, Geoffrey M. Voelker, Joe De-

Blasio, and Stefan Savage. Thanks to Mikhail Kolmogorov for his significant assistance, as well as translation help from Kirill Levchenko, Vector Guo Li, and Ivan Mikhailin. And thanks to Shawn Loveland, Elie Bursztein, Angelika Moscicki, Tadek Pietraszek, and Kashyap Puranik. This work was supported in part by NSF grants CNS-1629973 and CNS-1705050, and DHS grant AFRL-FA8750-18-2-0087.

Related articles on queue.acm.org

Criminal Code

Thomas Wadlow, Vlad Gorelik
<https://queue.acm.org/detail.cfm?id=1180192>

The Seven Deadly Sins of Linux Security

Bob Toxen
<https://queue.acm.org/detail.cfm?id=1255423>

The Web Won't Be Safe or Secure until We Break It

Jeremiah Grossman
<https://queue.acm.org/detail.cfm?id=2390758>

References

1. Anise, O., and Lady, K. State of the auth: Experiences and perceptions of multi-factor authentication. *Duo Security*, 2017; <https://duo.sc/2kmOBId>.
2. Cohen, W.W. Enron email dataset, 2015; <https://www.cs.cmu.edu/~enron/>.
3. Coonce, S. The most expensive lesson of my life: Details of SIM port hack, 2019; <http://bit.ly/2IGSD4Y>.
4. Google. Protect users with the Advanced Protection Program; <https://support.google.com/a/answer/9010419>.
5. Google. Protect your business with 2-Step Verification; <https://support.google.com/a/answer/175197>.
6. Google. Verify a user's identity with extra security; <https://support.google.com/a/answer/6002699>.
7. Honan, M. How Apple and Amazon security flaws led to my epic hacking. *Wired*; <https://www.wired.com/2012/08/apple-amazon-mat-honan-hacking/>.
8. Liu, S., Foster, I., Savage, S., Voelker, G.M., Saul, L.K. Who is .com? Learning to parse WHOIS records. In *Proceedings of the ACM Internet Measurement Conf.*, 2015, 369–380; <https://dl.acm.org/citation.cfm?id=2815675.2815693>.
9. Matishak, M. How Podesta became a cybersecurity poster child. *Politico* 2016; <https://politi.co/2m4fNmd>.
10. Mirian, A., DeBlasio, J., Savage, S., Voelker, G.M., Thomas, K. Hack for hire: Exploring the emerging market for account hijacking. In *Proceedings of the World Wide Web Conf.*, 2019, 1279–1289; <https://dl.acm.org/citation.cfm?id=3313489>.
11. Onaolapo, J., Mariconti, E., Stringhini, G. What happens after you are pwnd: Understanding the use of leaked webmail credentials in the wild. In *Proceedings of the ACM Internet Measurement Conf.*, 2016, 65–79; <https://dl.acm.org/citation.cfm?id=2987475>.
12. Thomas, K. et al. Framing dependencies introduced by underground commoditization. In *Proceedings of the Workshop on the Economics of Information Security*, 2015.
13. Thomas, K. et al. Data breaches, phishing, or malware? Understanding the risks of stolen credentials. In *Proceedings of the ACM Conf. Computer and Communications Security*, 2017, 1421–1434; <https://dl.acm.org/citation.cfm?id=3134067>.

Ariana Mirian is a Ph.D. student in the computer science and engineering department at the University of San Diego, CA, USA, where she focuses on understanding security and privacy via an empirical lens.

Copyright held by author/owner.
Publication rights licensed to ACM.