ACCOUNT TAKEOVER IN 2022

THE 24-BILLION PASSWORD PROBLEM

Digital Shadows Research Team

digital shadows_

ACCOUNT TAKEOVER IN 2022

EXECUTIVE SUMMARY

24,649,096,027 account usernames and passwords have been exposed by cyber-threat actors, as of this year. That's an immense number—one that should shake online account owners (all of us) to our very core. But despite this number, which grows bigger each year, and the deluge of reports highlighting the risk of insecure credentials, you still have a colleague—maybe more than one—who's carefully typing 123456 into a password field right now.

Credential abuse is something that happens to other people's accounts, right? Nope. It's everywhere, and your compromised passwords and usernames are enabling all kinds of threat actors to perform all kinds of account takeover (ATO) attacks. Your bank account might pay for a new TV purchase. Your colleagues might be persuaded to transfer money, after receiving emails supposedly from you. Your Twitter account might suddenly start spewing out spammy content. The malice varies, but the ATO risk consistently adds up to bad news. We all like hearing stories from the good old days about leaving your door unlocked at night, but doing the same with your accounts is inviting disaster.

Feeling any déjà vu? The Digital Shadows Photon Research team reported on ATO in 2020, revealing the scale of credentials available on cybercriminal locations (massive), and the ease with which actors were stealing, exploiting, and selling access to stolen accounts (extreme). Two years down the line, the situation is arguably no better…see the key findings of our research summarized below. Weak passwords abound, and ATO is interrupting services critical to every aspect of online life: working, streaming, ordering, paying, and just plain connecting. Raising security awareness of this topic can certainly help; but the ATO threat will remain endemic until the problems inherent to password use are resolved.

- **We collated more than 24 billion compromised credentials. That's a 65 percent increase from 2020, likely fueled by an enhanced ability to steal credentials through dedicated malware and social engineering, plus improved credential sharing.**
- **Within this data set, approximately 6.7 billion credentials had a unique username-and-password pairing, indicating that the credential combination was not duplicated across other databases: This was 1.7 billion more than found in 2020, highlighting the rate of compromise across completely new credential combinations.**
- **The most common password, 123456, represented 0.46 percent of the total of the 6.7 billion unique credentials. The top 100 most common passwords represented 2.77 percent of this number.**
- **Information-stealing malware persists as a significant threat to your credentials. Some of these tools can be bought for as little as $50, and some go for thousands, depending on functionality.**
- **Cybercriminal marketplaces and forums remain hot spots to buy and sell stolen credentials; we've alerted clients about exposed credentials advertisements 6.7 million times in the past 18 months. Several subscription services have also emerged, offering cybercriminals a premium service to purchase stolen credentials.**
- **The price of credentials depends on the account's age, the buyer's reputation, and the size of the data file on offer. Certain account types, like cryptocurrency-related accounts, also garner higher rates.**
- **Once credentials have been obtained, free, open-source credential-stuffing and password-cracking tools can give threat actors all the functionality required for a sophisticated attack to unlock passwords.**
- **Offline attacks usually produce the best results for cracking passwords; 49 of the top 50 most commonly used passwords could be cracked in less than a second. Adding a special character to a basic ten-character password adds about 90 minutes to that time. Adding two special characters boosts the offline cracking time to around 2 days and 4 hours.**
- **Financially motivated, state-sponsored, and ideologically motivated actors (hacktivists) have all used ATO as a conduit for their activity in 2022. This includes several attacks by the data extortionists known as Lapsus$ Group.**
- **Until passwordless authentication becomes mainstream, the best ways to minimize the likelihood and impact of ATO are simple controls and user education—use multi-factor authentication, password managers, and complex, unique passwords.**

WHAT'S NEW SINCE OUR LAST ATO OUTLOOK?

When our previous report was issued in July 2020, the COVID-19 pandemic was in full effect and working practices had changed radically. The shift to remote work has persisted, and WFH (working from home) is now firmly in the lexicon of every industry. The risks associated with remote services have dramatically increased as a result, and many organizations with insecure methods of authentication have become victims; the volume of ATO attacks has been skyrocketing since the start of the pandemic.

WFH isn't the only shift we've observed since our first ATO paper came out. Since the beginning of the Russia-Ukraine war in February 2022, the security landscape of the Western world has become increasingly unstable and is driving malicious cyber activity. The war has created a fertile environment for state-sponsored threat actors and cybercriminals to use wiper malware, perform distributed denial of service (DDoS), and deface websites, among other malicious cyber activities. The fallout has included several breaches of accounts and sensitive material associated with organizations in Ukraine and Russia.

Even in the face of these epic changes to the threat landscape, certain behaviors remain unchanged and open up opportunities for malicious activities. It's a well-known fact that basic cyber hygiene significantly lowers the risk of ATO, but many online users keep reusing passwords or creating vulnerable, easy-to-guess passwords. They're practically inviting attackers to compromise their accounts. This was recently demonstrated in Verizon's Data Breach Investigations Report (DBIR), which found that stolen credentials accounted for 50 percent of the 20,000 incidents analyzed by Verizon. This represented a 30 percent increase in use of stolen credentials found in the DBIR from 2017.

Figure 1 shows results of the Photon team's 2022 count of how many breached credentials are out there. 2019 seems to have been the standout year for the number of credentials found, but 2020 and 2021 both also returned approximately 5 billion new credentials each year. This tempo of breaching will probably continue in 2022, or increase, with the growing use of online services.

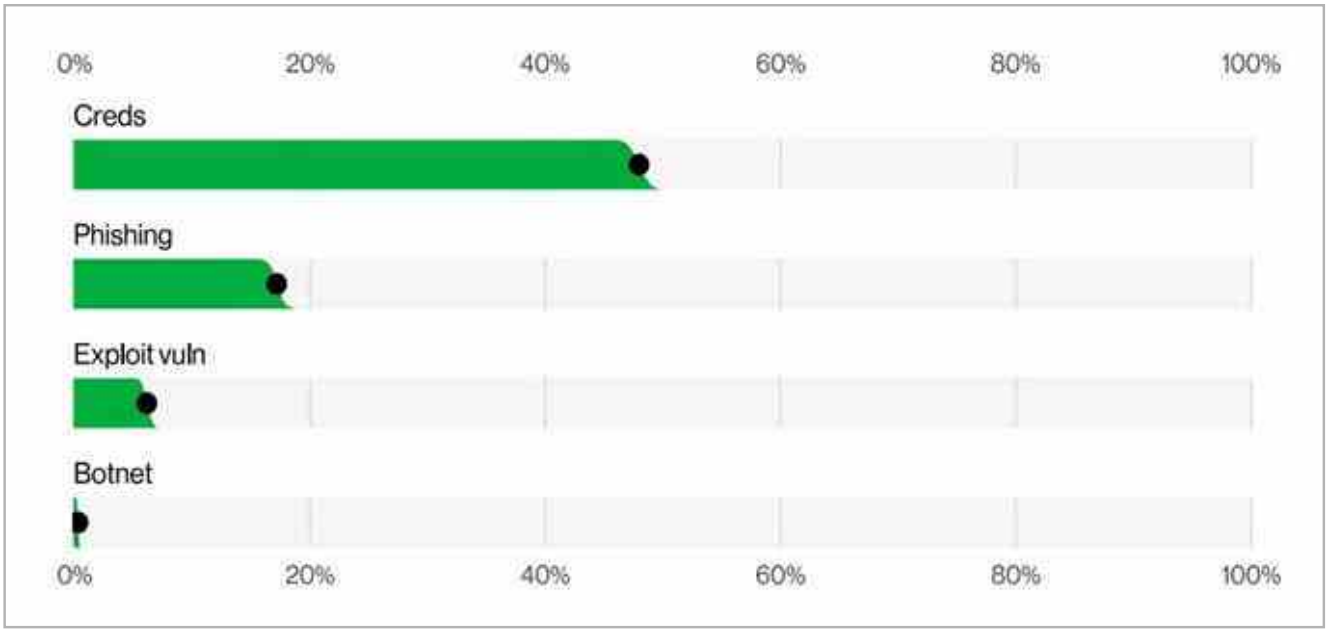

FIGURE 1 *Initial access vectors (Source: Verizon DBIR)*

FIGURE 2 *Number of credentials collated by Digital Shadows: 2016–21*

To address the pervasive password problems that lead to credential exposure and ATO, our society needs to move beyond passwords altogether. That's why tech giants Google, Apple, and Microsoft all recently announced their commitment to a "passwordless future". According to their stated objectives, in the coming years, our phones will store a FIDO (Fast Identity Online) credential, called a passkey, to unlock all our online accounts. We'll no longer need the plaintext passwords that are so easily compromised and exploited.

Passkeys can be considerably more secure than passwords because they're protected with cryptography. Also, they're only shown to your online account when you unlock your device. Although not a bulletproof solution, this method holds significant promise to reduce the number of ATO attacks.

IN THE COMING YEARS, OUR PHONES WILL STORE A FIDO (FAST IDENTITY ONLINE) CREDENTIAL, CALLED A PASSKEY, TO UNLOCK ALL OUR ONLINE ACCOUNTS. WE'LL NO LONGER NEED THE PLAINTEXT PASSWORDS THAT ARE SO EASILY COMPROMISED AND EXPLOITED.

3 WEAK SPOTS ENABLING THE ATO ATTACKER

THE EVER-EXPANDING DIGITAL FOOTPRINT

Modern business depends overwhelmingly on their digital presence, with a **large digital footprint** undergoing consistent change being difficult to manage. Think of a digital footprint as an expression of your company's entire online presence: its brand, reputation, and marketing and sales strategies. As your company grows, so does its digital footprint.

During change or other growth, mistakes happen...employees come and go, responsibilities change, assets are handed over or merged with different departments. Even if you're aware of that expanding footprint, it can be hard to stay on top of individual responsibilities or understand exactly what needs to be secured, not to mention gain visibility of all assets so you can spot any misuse.

Your company's digital footprint is invariably linked to its attack surface. Think of that as the total number of entry points where an unauthorized user can access a company system and extract data. The smaller the attack surface, the easier it is to protect. But that's often easier said than done. A lack of understanding of what assets you own, what materials they constitute, or how to remediate any problems are key areas where Digital Shadows assist clients.

THE AUTHENTICATION BLIND SPOT

ATO is also helped by lack of consistent **authentication** across accounts. Passwords are a headache—we all know this—and users are still bypassing the abundant password managers and other solutions, using unsafe methods to log in to their accounts. Nobody wants to remember a 16-character alphanumeric-symbol password. It's just more convenient to use an old tried-and-tested password for multiple accounts, or save it in a browser or other unsafe location.

Many people believe that despite the risks, their credentials won't be cracked. It's always going to be someone else, affecting a service they don't use. Well, cybercriminals are exploiting that ignorance, easily identifying, harvesting, and cracking credentials using modern cracking tools.

THE TOO-LATE ATTEMPTS AT ACCOUNT PROTECTION

The combination of increasing attack surfaces and unsafe authentication is leading to **more accounts being exposed than ever before**. Not a week goes by without a data breach affecting a popular service provider, leaving users scrambling to change their credentials or otherwise minimize the risk.

Once credentials have been compromised, their sale often attracts buyers with considerable enthusiasm. Files containing breached credentials are either sold at a premium, for the access and opportunities they provide buyers with, or given away for a very low price, or for free, granting a wide range of actors the chance to make use of them. Credential misuse is one of the primary enablers of several forms of cybercrime.

EASY AS 123: THE LIFECYCLE OF AN ATO ATTACK

As with any cyber attack, ATO starts with a mistake, a misconfiguration, or another oversight that provides an opportunity to someone with malicious intent. It's often difficult to spot before it's too late.

There are many scenarios where ATO can flourish, but a typical lifecycle involves identifying a susceptible service or user, attempting to acquire accounts, verifying whether they can be used across other services, and exploiting these accounts for nefarious purposes.

WITH PASSWORDS USED ACROSS MULTIPLE ACCOUNTS, EVEN A BREACH OF A RELATIVELY BENIGN ACCOUNT CAN TURN INTO SOMETHING MUCH MORE SERIOUS; THREAT ACTORS CAN EASILY PIVOT TO OTHER ACCOUNTS.

STAGE 1: RECONNAISSANCE

The intrepid ATO attacker's first priority is reconnaissance, to identify susceptible accounts. They might look at previous data breaches to notice which organizations are frequent victims. Or they could zero in on specific individuals or organizations whose associated accounts seem to offer value. And that pernicious digital footprint can also reveal a wealth of information the attacker can use to craft their attack. (That's something Digital Shadows can help you with!)

Having nailed down their target, the attacker starts mapping login portals that use username-password combinations. This might be an online banking portal, a social media login page, or a login platform to process your tax returns. By enumerating the type of technology used on the site and any subdomains, they can determine which might be susceptible to exploits or credential-based attacks. In essence, it's an information-gathering exercise, aimed at creating a full inventory of all Internet-connected devices and domain names of a target company.

There's an absolute glut of tools that can be used for recon, and a whole community of cybercriminals who'll be happy to assist—often for a small fee, of course (there's probably an analogy for teaching a cybercriminal to phish in here, somewhere...). We found several posts on the popular Russian-language cybercriminal forum Exploit that illustrated a daily exchange of information among collaborators.

Attackers often have a specific target in mind; in Figure 3 you can see a request for assistance in buying accounts associated with French homeowners. (Side note: These requests typically reflect the capability of the attacker to exploit accounts associated with a certain geographic region; in this case, the requester was probably French.)

Figure 3 shows a request for recommendations of software that can scan the code of certain URLs, to indicate "the presence of certain links or scripts". In this case, the attacker was probably looking to identify website vulnerabilities to exploit for credential stuffing.

Fierce is a particularly useful tool for identifying susceptible services. It's a lightweight scanner that helps locate non-contiguous IP space and hostnames on specified domains. (If you're unaware of non-contiguous IP space, you're not alone. It refers to IP addresses that can't be summarized.)

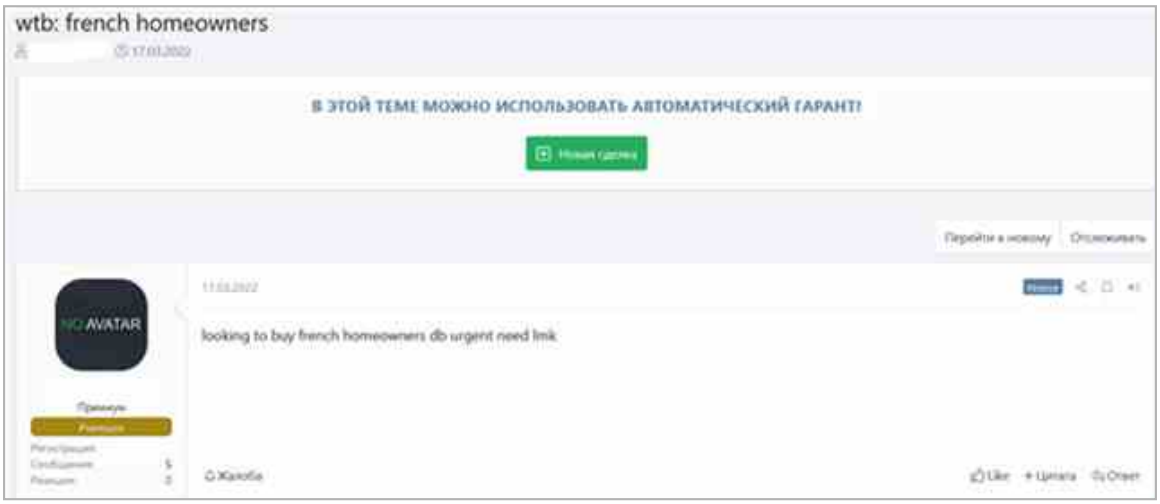

FIGURE 3 *Cybercriminal requesting accounts of French homeowners on Exploit*

FIGURE 4 *Cybercriminal user requesting exploits for website enumeration on Exploit*

Fierce was designed as a precursor to other enumeration tools that require users to already know what IP space they're looking at. Fierce doesn't perform exploitation and doesn't scan the whole Internet indiscriminately. It's a specialist, aimed at specifically locating likely targets inside and outside a corporate network. This can be particularly useful to a malicious actor attempting to identify misconfigured networks that leak internal address space.

Fierce can be used to direct malware at blind spots in a network, or, as in our example, target inactive web portals that might have been left online. This goes back to our earlier point about the dangers of a large and unmanaged attack surface; more entryways to your house demand more effort to close them before you head out. The same applies to your network.

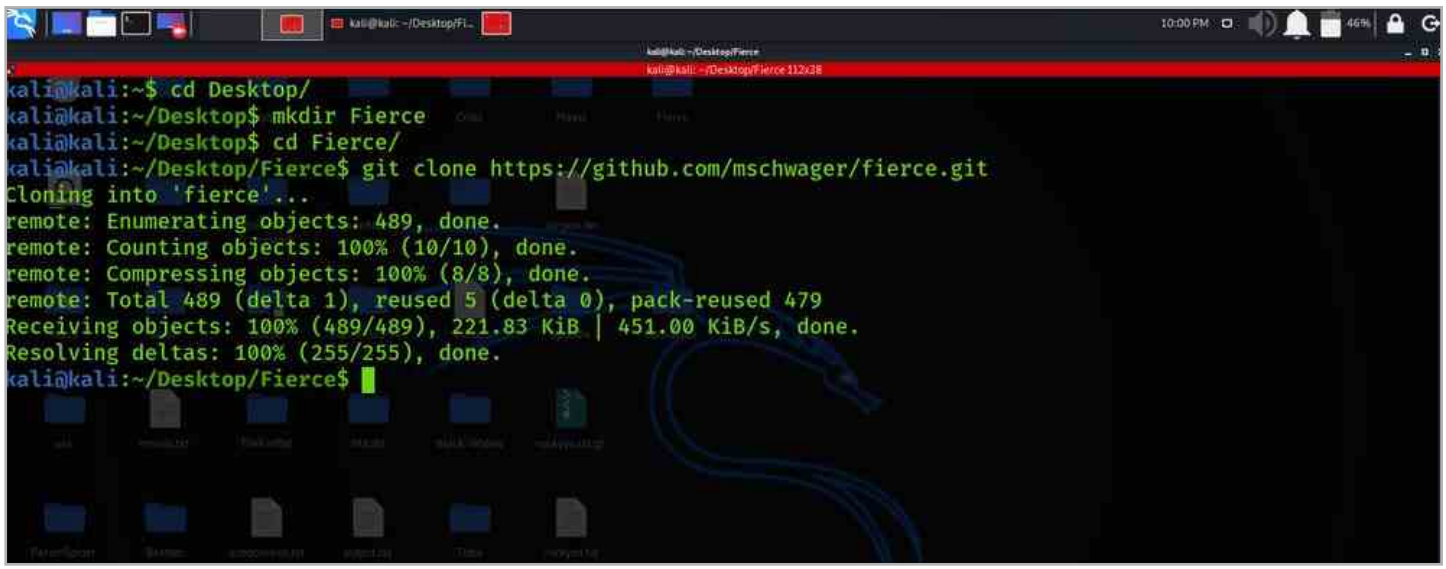

FIGURE 5 *The Fierce DNS reconnaissance tool living up to its name*

STAGE 2: ACQUISITION

The target has been identified—whether it's susceptible accounts or a susceptible organization hosting accounts. Time for our intrepid ATO attacker to move forward. They can choose from a wealth of methods to steal the accounts' credentials from their rightful owners. Methods constantly evolve, so alongside the proven approaches frequently surface new and creative methods. (Hard truth ahead: As with anything in cyber security, you can improve your overall resilience but threat actors will always adapt, refresh their approach, and come back later. All we can do is roll with the punches and make sure the basics are done right, minimizing mistakes as much as possible.)

One of the most common credential-stealing methods is using simple **phishing emails or other social-engineering techniques,** artfully tricking users into actually handing over their credentials or brutally harvesting credentials with malware. We've issued grave warnings about the threats of social engineering, and how phishing, in particular, is still a major problem, even in 2022. Because, ultimately, social engineering works: Humans are flawed, inquisitive, and curious…it's easy to be taken in with the right lure.

Email ATO is a top way to harvest your credentials. The threat actor sends phishing emails to predetermined email addresses that the attacker already is aware of, but for which they do not have a password. You, as a message recipient, are redirected to malicious infrastructure that the attackers own—picture a webpage that mimics your favorite entertainment streaming website. Hook, line, and sinker, you unwittingly enter your credentials, giving the attacker what they need to conduct a more complex attack.

A phishing email can also be a conduit to distribute **credential-stealing malware**—shout out to our 2020 report's number-one way to acquire credentials. Banking trojans, in particular, can be a highly lucrative option. You've probably heard of the likes of "Emotet", "Trickbot", and "Ursnif", but a new generation of banking trojans is also enabling cybercriminals to harvest financial personally identifiable information (PII).

A lot of credential-stealing malware is free and being openly advertised in open areas of criminal forums. On the Russian-language cybercriminal forum XSS, one user shared a consolidated list of free and available malware being hosted on GitHub, just waiting for a motivated malcontent to try out against a live target.

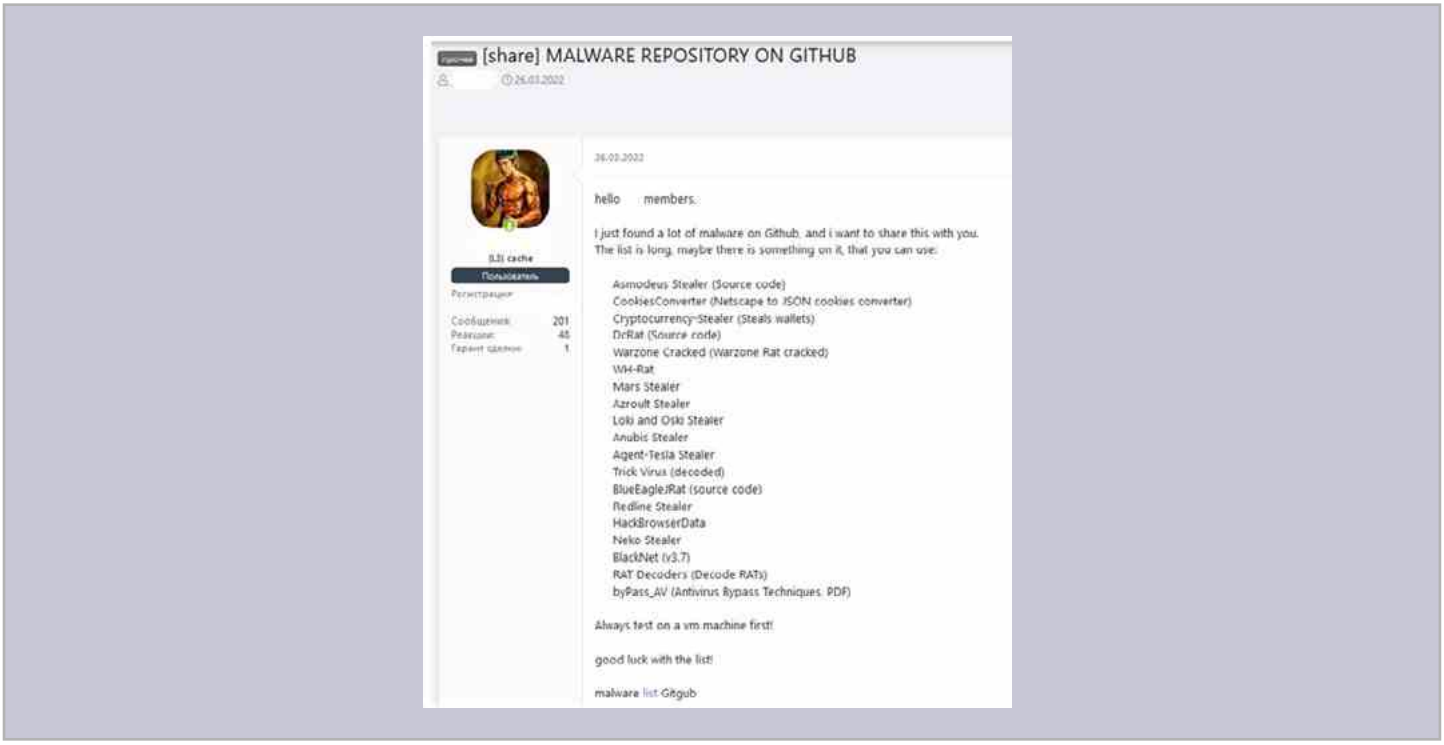

FIGURE 6 *List of credential-stealing malware stored on GitHub, shared on XSS*

Let's talk through some of the most popular credential-acquisition helpers that cybercriminals are relying on.

Redline information stealer

Time and time again, through the course of our research—and almost weekly while assisting clients with incidents—we've seen the Redline information stealer (infostealer) crop up. This commonly used, commodity malware targets most popular web browsers, including Chrome, Edge, Opera, and Firefox. It is sold for roughly $200 on cybercriminal forums and can be deployed without much technical know-how.

FIGURE 7 *Redline official Telegram channel*

Redline boasts more than a few capabilities. But key to ATO attackers is information collection—including that saved to browsers, like credentials, credit card details, and cookies, plus system information about software and hardware used with the infected device. It can detect processes and anti-malware software on the system, which helps the attacker fine-tune their approach.

Redline is a big problem for corporate accounts, not just individuals. The latter can offer quick monetization of stolen accounts and financial information, but corporate accounts can aid nefarious attackers looking for an easy way into a particular network. So although Redline is most commonly associated with cybercriminals, it's also feasibly useful to a nation-state or non-financially motivated group. We've seen this malware in attacks by several groups, including the notorious Lapsus$ Group data extortionists.

How is Redline spread? Well, you've probably guessed it: a combination of phishing and other pretty basic methods. Several Redline campaigns in late 2021 used phishing lures, like fake advertising requests, holiday gift guides, and website promotions. Recent phish bait includes fake deals associated with cryptocurrency company Binance's non-fungible token (NFT) mystery box.Users receive a randomized NFT and happily interact with these boxes, hoping they'll receive a unique or rare item at a bargain price. In that campaign, threat actors even created YouTube videos to add an air of legitimacy to the scam.

These scams often deliver malicious Microsoft XLL (an Excel add-in extension) files that result in Redline installation. In some cases, fake infrastructure, including spoofed websites, is used to host the XLL files. Abuse of XLL is fairly common, given that the add-in enables developers to extend the functionality of Excel by reading and writing data, importing data from other sources, or creating custom functions to perform various tasks. Of course, this represents a huge opportunity for malicious actors. If this kind of file ever lands in your inbox from an unexpected source, it goes without saying that you should treat it with extreme suspicion.

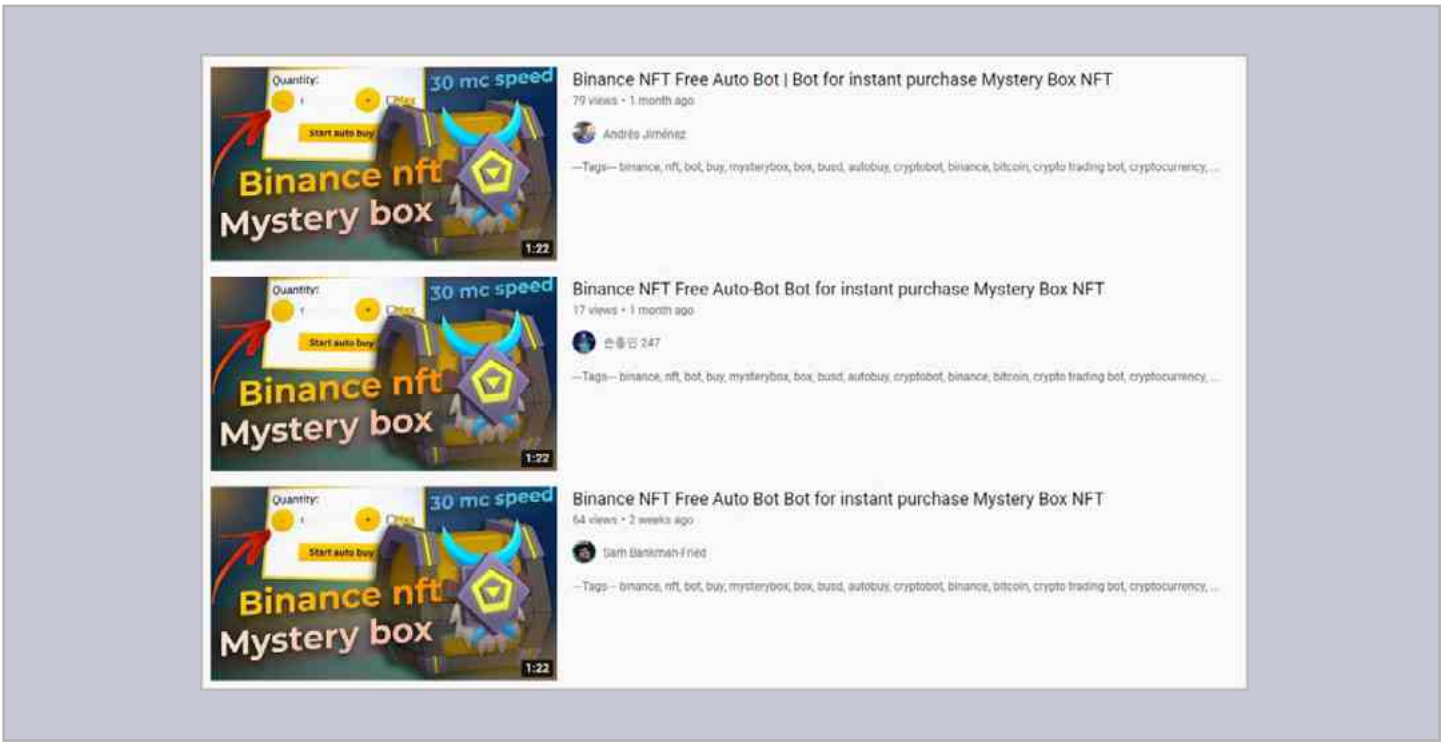

FIGURE 8 *YouTube videos facilitating NFT box scam (Source: NetSkope)*

Once Redline has infected a system, an additional task runs in the background to collect all the system information it can (in Figure 8, for example, the task is named AddinProcess.exe). This is how account credentials can be collected, especially if they're not stored safely (e.g. they're stored in a web browser).

Redline's success manifests through a series of failures: It's spread through interaction with a suspicious email that isn't blocked by an email gateway, and credentials that are unsafely saved are collated. It really is that simple.

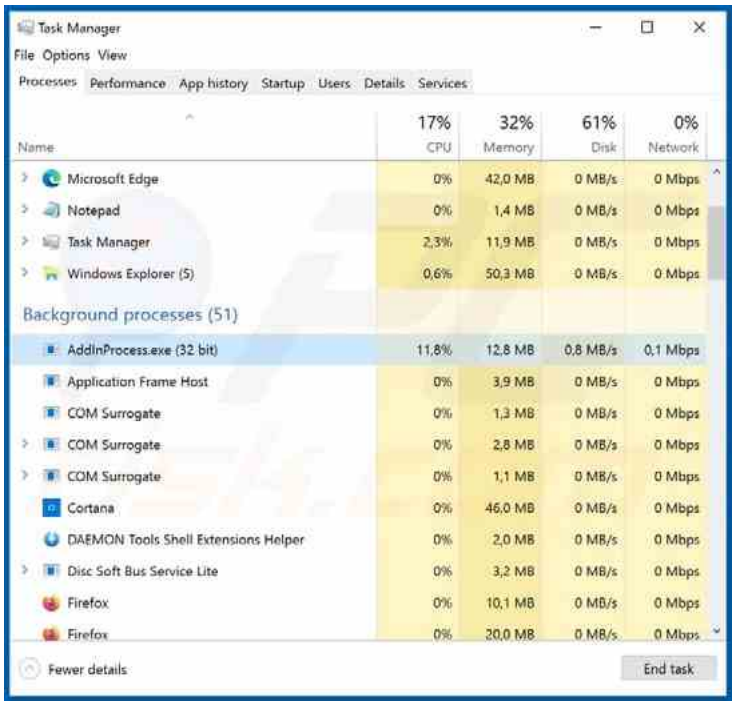

FIGURE 9 *Redline malware process running in background (Source: PCRisk)*

Automated vending carts

Redline and several other commodity infostealers have been used to great effect in boosting automated vending carts (AVC) service development and expansion. We've warned about AVCs before, highlighting the dangers of these websites that enable the sale of illicitly obtained goods without the need for buyer-vendor interaction. AVCs allow quick and seamless transactions for a wide range of items—financial information, app accounts, email account credentials, credit card details, bank account data, stolen log data, fingerprints, remote desktop protocol (RDP) access…the list goes on. We monitor several AVCs, including Genesis, 2Easy, and Russian Market.

AVCs are longstanding hot spots for cybercriminals, and have become even more popular since our last report. Russian Market isn't exactly the cybercriminal eBay, but it's not far off, and the same goes for other AVCs. At a very high level, they dramatically lower the technical-knowledge threshold for threat actors to enter the cybercriminal world.

A highly customizable interface enables cybercriminals to filter by factors like geography, IP address range, or level of access. And low prices make purchasing even easier: A credit-card listing can be bought for a mere $12–30. Overall, AVCs are highly effective for obtaining stolen accounts, and much AVC activity is facilitated through commodity malware, like Redline.

Other platforms to purchase or rent

Beyond social engineering, malware, and AVCs, maybe the most common method to gain hold of credentials is buying them through a dedicated cybercriminal marketplace or forum. These are a critical element of the cybercrime ecosystem, permitting a range of malicious activity. The credentials they advertise have been compromised via any number of the methods we described above…exfiltrated from previous access, taken through stealer malware, resulting from "log parsing"[1], or handed over by socially engineered victims.

Credentials are most commonly sold in marketplaces, or through forums' commercial sections. In Figure 10, you can see how a threat actor advertised 132 raw logs that were stolen using Redline. Whoever buys these logs can parse the results and potentially extract any credentials that the affected account's owner may have entered.

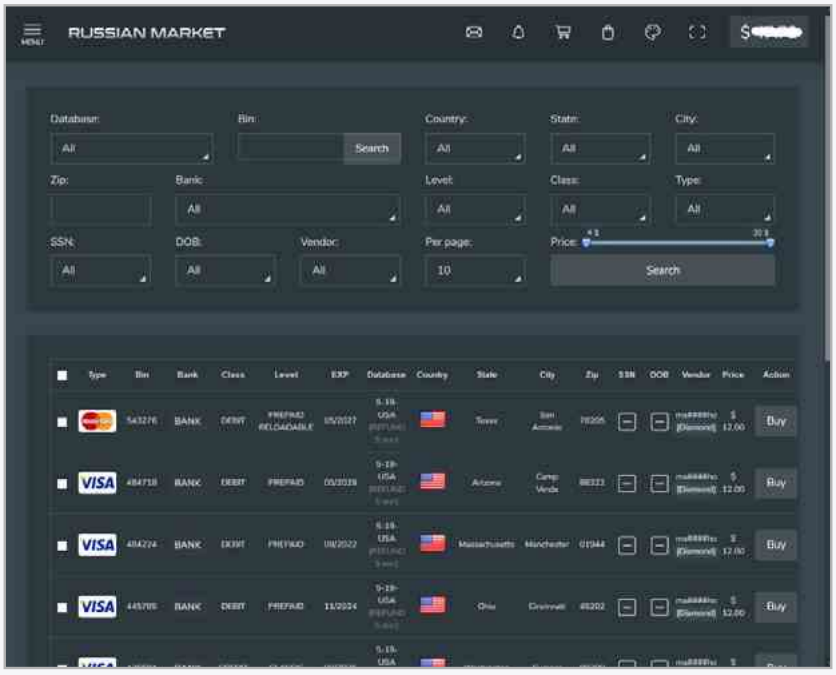

FIGURE 10 *Russian Market—a walking, talking, eBay for cybercrime*

[1] Splitting data into chunks of information that are easy to manipulate and store; for cybercriminals, log parsing involves identifying any credentials stored within the logs' data.

FIGURE 11 *Cybercriminal selling logs stolen through the Redline infostealer*

For cybercriminals who don't feel like visiting forums or marketplaces, there are also dedicated services they can subscribe to for a monthly fee. Why pay? Well, each database listed in such a service has a limited "shelf life": The usefulness of the accounts reduces fairly quickly, as more users gain access and the owners are potentially alerted to fraudulent activity; credentials can be changed and access lost. So there's value in forking out money to be informed of account listings, before other actors get their hands on them.

BRIKK is a good example of such a service. It's an English-language subscription service that advertises compromised email-password and username-password credential pairs that are intended for use in password-cracking attacks. BRIKK offers subscriptions on a weekly, monthly, or lifetime basis, with the weekly subscriptions priced at $222. Users must purchase a subscription to view any data hosted on the BRIKK site, which has earned many positive reviews on its dedicated Discord server and multiple cybercriminal forums. BRIKK demonstrates something we allude to a lot in our research: the growing professionalization of criminal services. If you're wondering what the difference between a cloud and crack subscription, crack refers to a subscription permitting access to dedicated cracking tools also available on BRIKK.

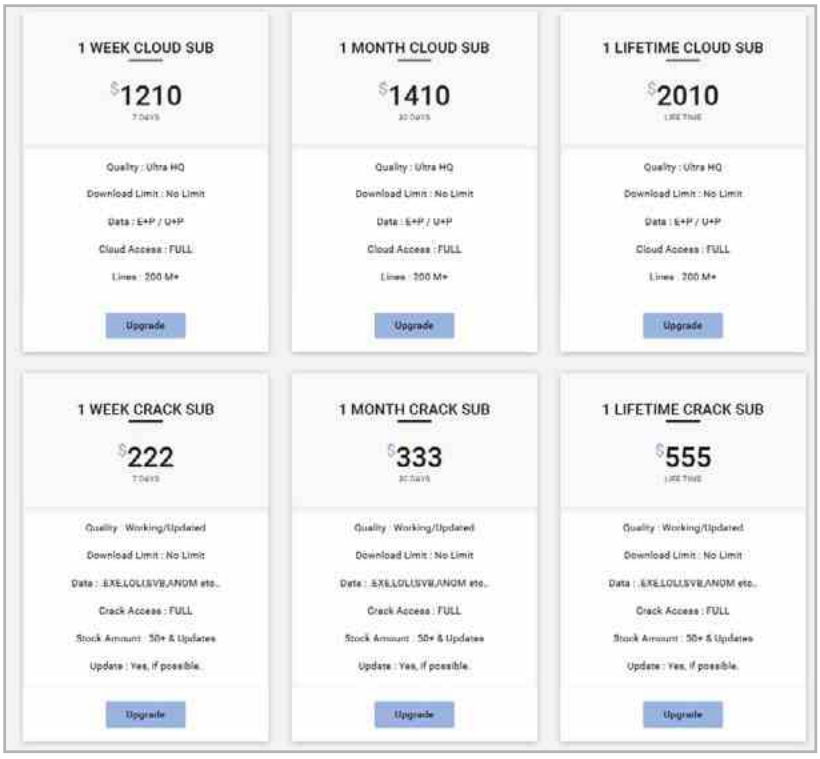

FIGURE 12 *Subscription models available for BRIKK*

THE COST OF CREDENTIALS: WHAT'S BEHIND THE PRICE TAG?

What drives the pricing of account credentials? There are a few factors at play here. The first is the validity rate: the number of credential pairs that are legitimate and working at the time of a sale. Perception of validity is enhanced by a seller's history (i.e. they've made several sales in the past that turned out to be valid). Newer or less-experienced sellers usually don't warrant high prices for credentials.

Other factors are the size of the data file being sold, and whether the passwords are plaintext or encrypted; if they're stored in plaintext or an easy-to-crack algorithm, they'll typically be more expensive. Passwords that are hashed and salted (a unique, random string of characters known only to the original service has been added) tend to be cheaper, as they're more difficult for the buyer to crack.

Recency is also a factor—being the first person to use a data set can mean gaining access before an account owner has had a chance to change their credentials or take other mitigation steps. Certain types of accounts offer other enticing opportunities. We've seen a glut of forum posts from cybercriminals specifically seeking cryptocurrency accounts. These are coveted because money stolen from them can be moved quickly, with few options for the victim to restore any lost funds.

Figure 13 shows a forum post about a database of registrant information of 312,000 CoinMarketCap users; even though it's not specifically related to cryptocurrency wallets, this data could be calibrated to catalyze social engineering of the account holders.

Another way sellers make a credential data file more appealing, and merit a higher price, is including cookies/browser extensions and proxies (we elaborate on this in the next section). These savvy salespeople know that their buyers will probably push their newly acquired accounts through a credential stuffing tool once purchased, to verify the accounts and conduct ATO en masse.

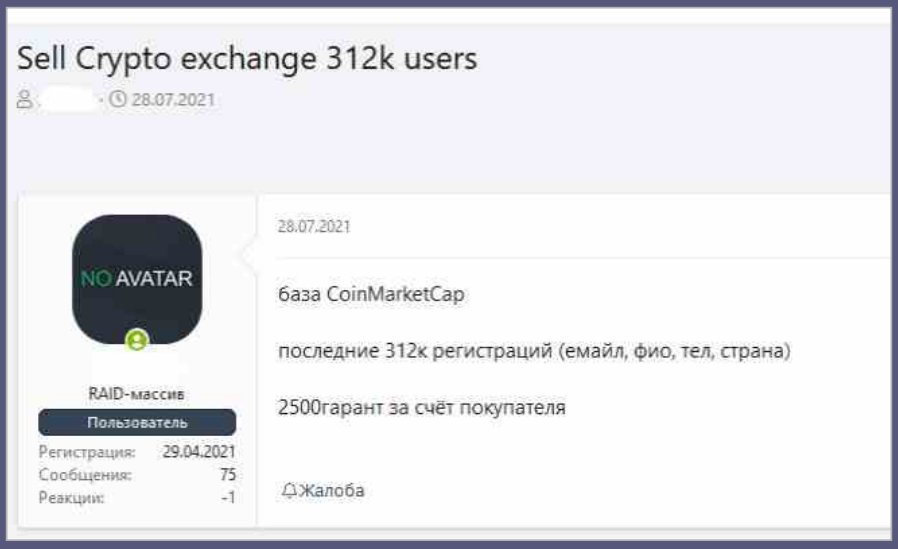

FIGURE 13 *XSS user selling CoinMarketCap database for $2500*

STAGE 3: VERIFY

And so our intrepid attacker finds themself in the third stage of the attack lifecycle. It's another fork in the road, with viable options in many directions that eventually lead to the same destination: ATO.

The incredible ease of credential stuffing

After the attacker has obtained credentials, they can use credential stuffing to consolidate them and facilitate ATO for other services across the Internet. In case you didn't pick up on this earlier: We're not talking about a handful of credentials; often the attacker potentially has thousands, or even tens of thousands, at their disposal. Attacks are simple and the risk is real.

Credential stuffing is a major problem because of weak passwords, insufficient controls, and the systemic re-use of passwords across multiple accounts. But who has the memory to recall a single unique password for each of hundreds of services? Out of convenience, many individuals use a simple password that can be easily remembered, either something pertinent to their own life, or staring at them on the keyboard. The risk this poses was demonstrated when General Motors, disclosed a credential stuffing attack on 23 May 2022. The attack resulted in an exposure of customers' data, allowing actors to redeem stolen customer reward points for gift cards.

The lifecycle of a credential stuffing attack can vary—and it is, in general, cyclical—but typically follows several defined steps, as shown in Figure 13.

A GOOD WAY TO THINK OF CREDENTIAL STUFFING IS OBTAINING A BAG FULL OF KEYS, AND TRYING TO UNLOCK A SERIES OF DOORS. THESE DOORS REPRESENT THE SITES AND SERVICES YOU USE EVERY DAY; THEY MIGHT OPEN UP YOUR SOCIAL MEDIA ACCOUNTS, YOUR EMPLOYER'S EXTERNAL PORTAL LOGIN, AND—PROBABLY MOST WORRYING—YOUR BANK ACCOUNT.

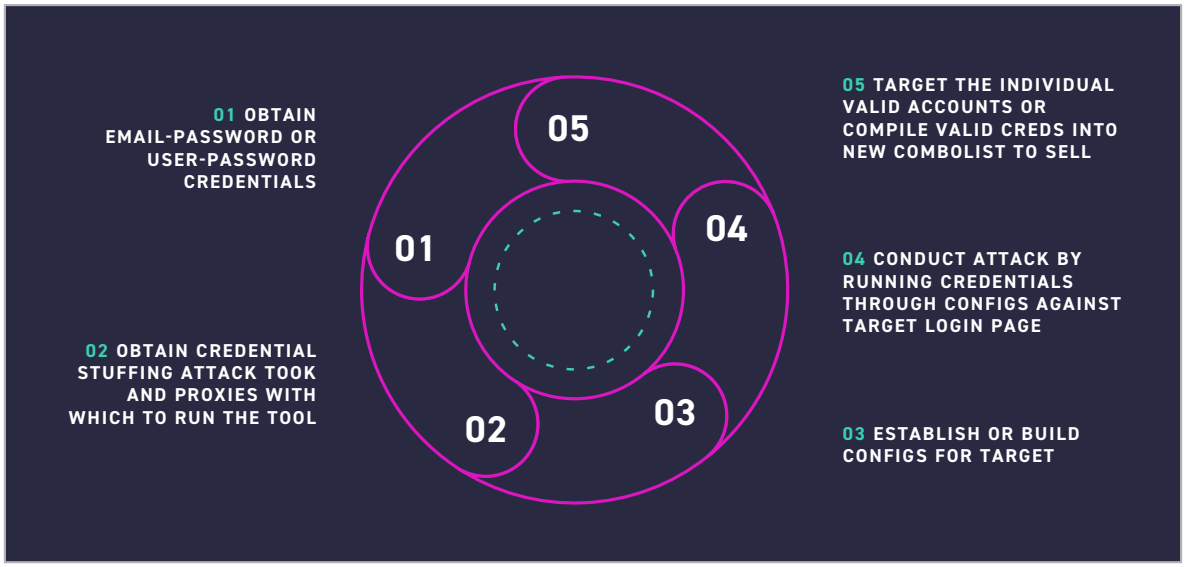

FIGURE 14 *Typical lifecycle of a credential stuffing attack*

It's time to talk about how immensely easy it is to perform credential stuffing, with readily available tools and very few skills required. (This may hurt.)

OpenBullet

OpenBullet is the most popular choice for credential stuffing. It's a freely available cross-platform automation suite, powered by .NET,[2] that enables the user to perform requests toward a target web app and offers an array of tools to work with the results. OpenBullet was originally released in April 2019 on GitHub, as a penetration testing tool intended for security researchers. What's that saying about the road to hell and good intentions?

This tool is one of several available to cybercriminals for DDoS or credential stuffing attacks on a website. And it's easy: Just input the URL you want to attack, load in the relevant config (more on that down below), and add the list of credentials you've managed to seize. OpenBullet is a favorite with average cybercriminals, because it's free to download and use, routinely updated, simple to use, and has an active community dedicated to its support. It's enabling a new generation of underdogs—low-skill/knowledge hackers—to disrupt organizations.

You'll find OpenBullet on GitHub, but on the main page the developer highlights, with remarkable prescience, that any illegal use of the tool is the responsibility of the user.

Once OpenBullet's relevant ZIP files are downloaded and installed on the system, the tool can be executed and the user can navigate to the home dashboard via the browser. As with all credential stuffing tools, use of OpenBullet is facilitated through the use of several precursor items of information, notably proxies, configs, and wordlists. These can be facilitated through an easy-to-use interface. Several dedicated tabs allow a cybercriminal to manage and monitor current jobs, update configs and proxies, and accumulate successful hits.

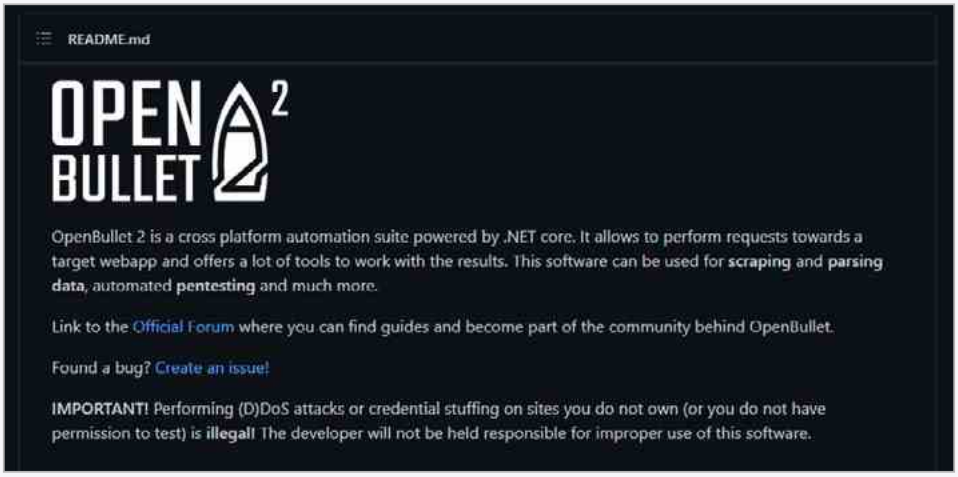

FIGURE 15 *OpenBullet system info*

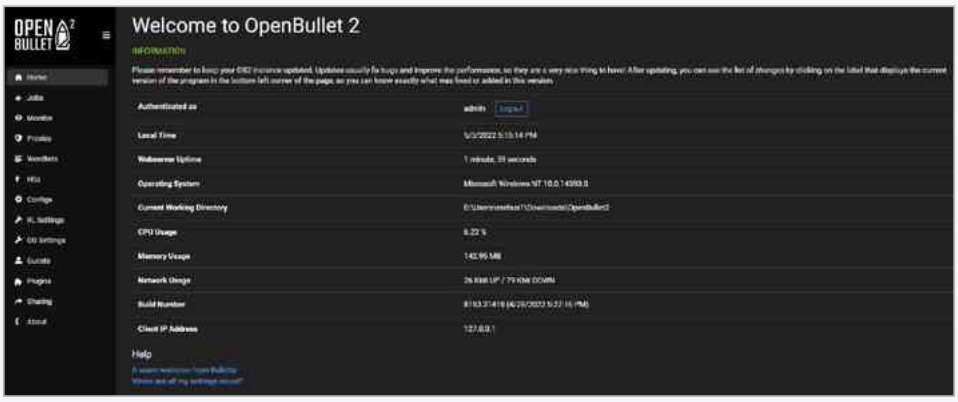

FIGURE 16 *OpenBullet homepage*

[2] .NET is a free, open-source, managed computer software framework for Windows, Linux, and macOS operating systems.

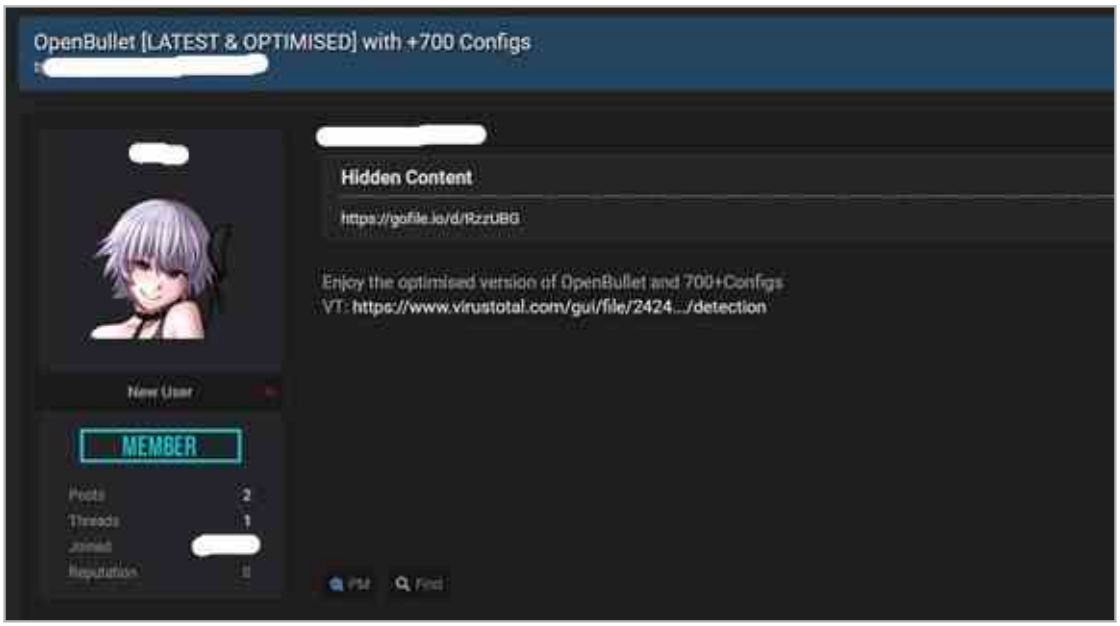

FIGURE 17 *RaidForums user advertising Open Bullet configs*

Configs

Let's talk configs, which are needed for use of OpenBullet and other platforms. Config files contain instructions for automating certain actions, usually for web resources. They include metadata (like a name, an author, and an icon) that acts as a README file showing how to use it or what it's for, some settings, and instructions in either the custom LoliCode scripting language or pure C#. In essence, a config tells the credential stuffing tool where to direct its usernames and passwords, and how to determine whether OpenBullet has successfully logged in or not. If session cookies are also required by a service, they would be used to ascertain that the login session has not been to a new device.

A config file can be created by the user for a target, or downloaded (sometimes for free) from certain cybercriminal platforms. Once the file has been acquir8d, the user can simply upload it to OpenBullet or whatever tool needs it, and provide some simple customization identifiers. Configs are routinely advertised for sale on cybercriminal forums, typically specifying the sites and services they can be targeted against.

After a config is uploaded to OpenBullet, the user can retrieve a wordlist from across the clear and dark web to be used by the tool for the attack process. Wordlists are essentially a list of common passwords that an actor can use to attempt password cracking, which are extensively available in cybercriminal spaces.

OpenBullet also has a range of advanced controls for setting appropriate cookies or changing proxy settings that might be required to pass through the login page of a given target (see the section that follows this one). Once a "job" is initiated, OpenBullet fires thousands of username-password combinations to the desired website, reporting any success within seconds. Once there's a match, the attacker can use the winning credentials manually, to crack into accounts to commit online fraud, or sell for profit.

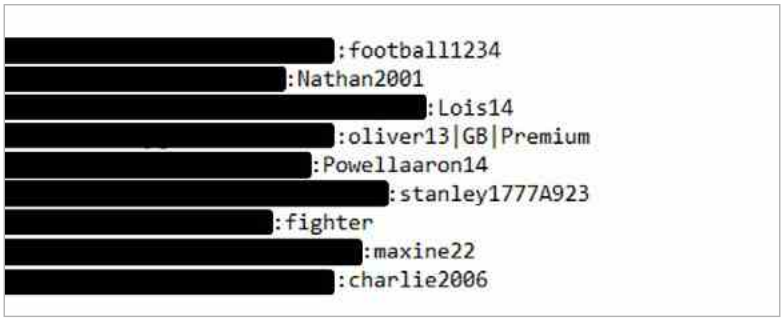

FIGURE 14 *A simple wordlist found online (email identifiers redacted)*

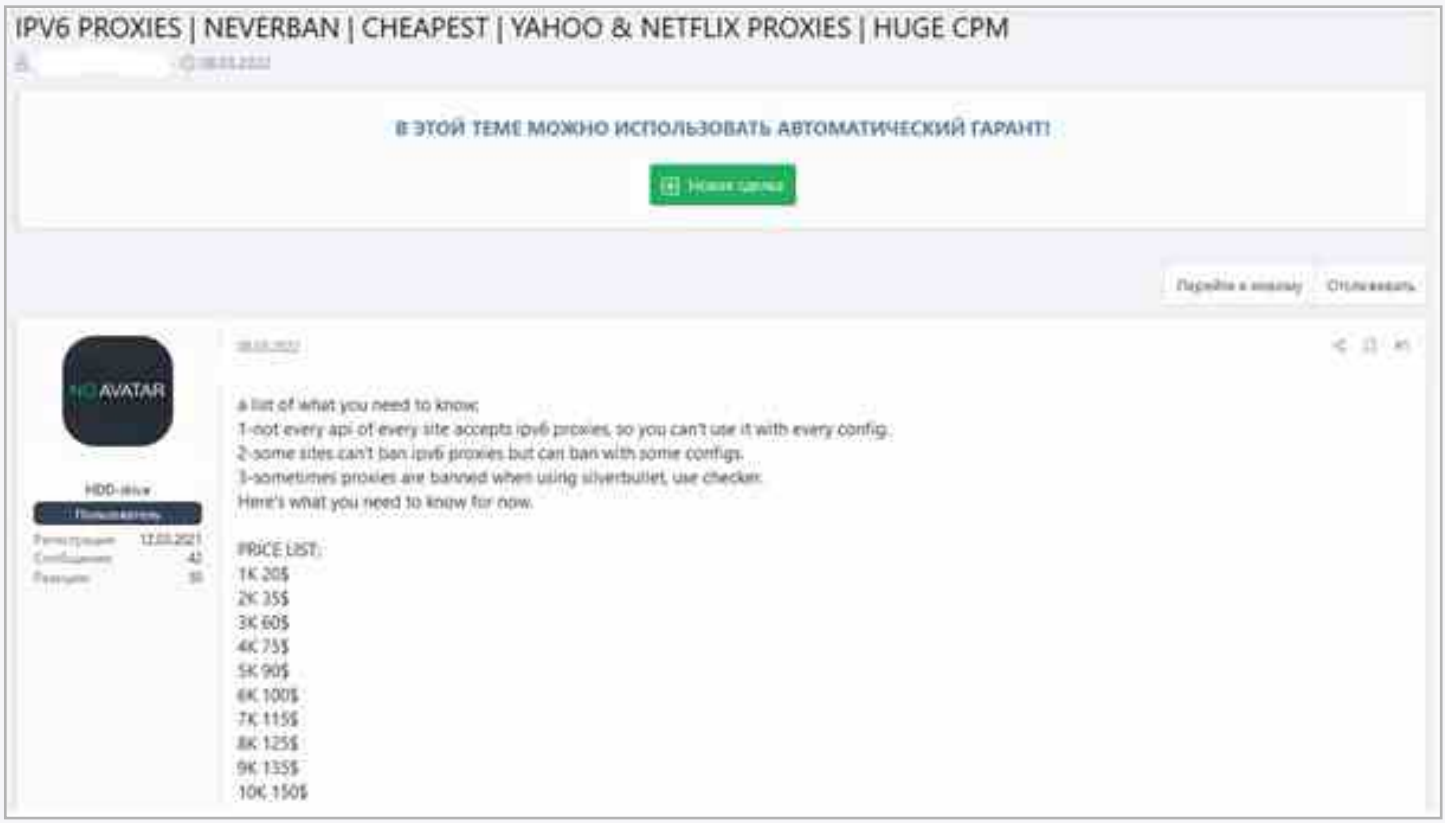

FIGURE 19 *IP version 6 (IPv6) proxies advertised for sale on XSS*

Proxies

Proxies are also key to successful credential stuffing attacks. Proxies allow cybercriminals to make multiple login attempts using a different IP address for each attempt. What's more, they can set up the time between each attempt, to avoid alarms triggered by unusual login activity (loads of attempts in a very short time); without those helpful proxies dynamically changing IP addresses for each request, the attack would be locked out fairly quickly.

PROXIES ARE OFTEN SOLD ALONG WITH CREDENTIALS—BY SELLERS WHO FULLY RECOGNIZE THE FUTURE USE OF WHAT THEY'RE SELLING. IF SOLD ON THEIR OWN, PROXY PRICES RANGE FROM ABOUT USD 20–150 FOR 1,000–10,000 PROXIES.

A DIRECT APPROACH: PASSWORD CRACKING

Not into credential stuffing? Not a problem. A dedicated password cracking tool can also be used to crack—i.e. verify—passwords. Instead of using plaintext and unencrypted passwords to attempt logins on various sites and services, this tool instead attempts to crack passwords that have been encrypted into a hash. (Hashing is a typically a one-way process to make plaintext passwords something that should be computationally indecipherable.) Selling hashed passwords to other threat actors is something the Photon team sees a lot, as well as requests for techniques to crack hashed passwords.

Cracking passwords is a process separate from credential stuffing, but it's often used in the same attack as a credential stuffing tool, potentially to provide an answer to the missing half of the credential pairing that the threat actor otherwise may not have. Once this puzzle has been solved, the attacker can turn to their credential stuffing tool and hugely expand the scope of their campaign.

The techniques used by credential stuffers and password crackers differ, but both work because—here we go again—most passwords in use today are inherently weak and predictable. If your password is longer and more complex, it stands a much better chance if someone tries to reverse the hash.

Online and offline attacks

Password cracking is achieved online, using a live service, or offline, if the attacker already has access to stored hashes and wants to crack the password on their own system. The difference between these two types of password cracking attempts is major.

An offline attack enables the attacker to test out a bunch of cracking processes without the fear of sounding any alarms or triggering account lockouts. This process is invisible to security teams and can be done in the attacker's own time, for as long as it takes, without fear of detection. They crack a password, identify the right combination, and log in to the application in a single instance. Network speeds of the service they are attempting to log in to aren't important; they're limited only by the speed of the computer doing the cracking.

An online, or live, password cracking attack is much more complicated. Each username-password combination must be sent over the network to an authentication server and then the respective server will respond accordingly. The time for this process depends on the application server's speed, and the speed of the network, but usually only three to five login attempts are allowed each second. This makes many online password cracking attempts much more difficult, and runs the risk of alerting defenders to malicious activity.

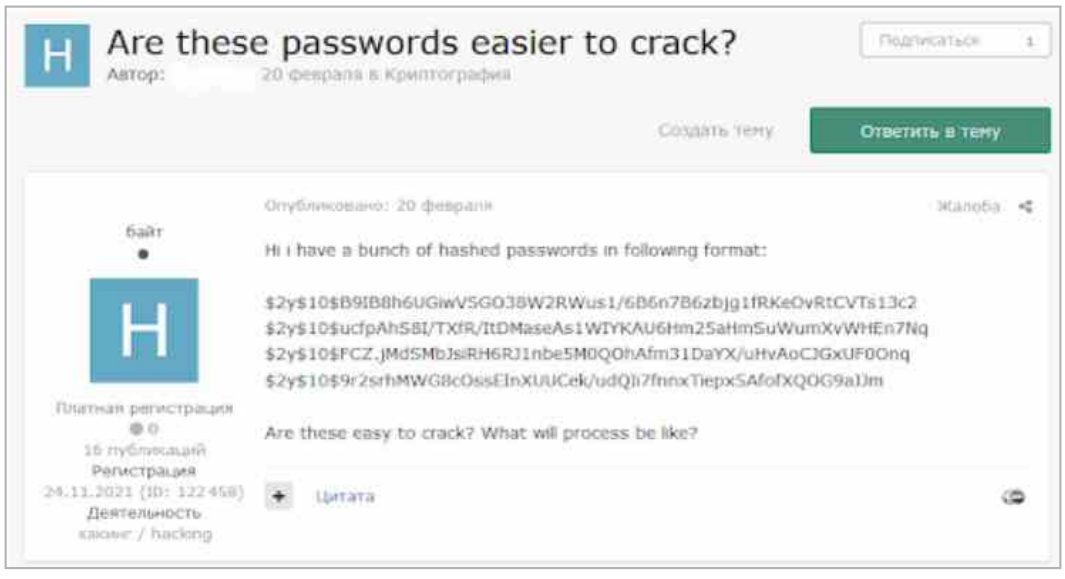

FIGURE 20 *Cybercriminal forum user seeks password cracking techniques*

Hashcat cracking tool

In 2022, cybercriminals—and also defenders—have been making great use of the offline cracking tool Hashcat. It's a series of free, open-source, readily available password crackers that offer the fast, accurate unearthing of a password from an encrypted hash, relatively simply. Its value to cybercriminals is obvious, but Hashcat can also enable your system administrator to ascertain the strength of passwords being used on your network. If you'd like to get your hands on Hashcat, it comes with many open-source penetration testing frameworks, including Kali Linux.

Basically, Hashcat works like this: Guess a password, hash it, compare the result to the original hashed password. If these passwords are the same, success! If not, Hashcat purrs along until it finds the right answer. This process can be astonishingly quick or agonizingly slow, depending on a password's complexity. Hashcat reportedly can uncover passwords encrypted in several formats, including Microsoft's LM hash algorithm, MD4, MD5, the SHA hash family, and the Unix crypt format.

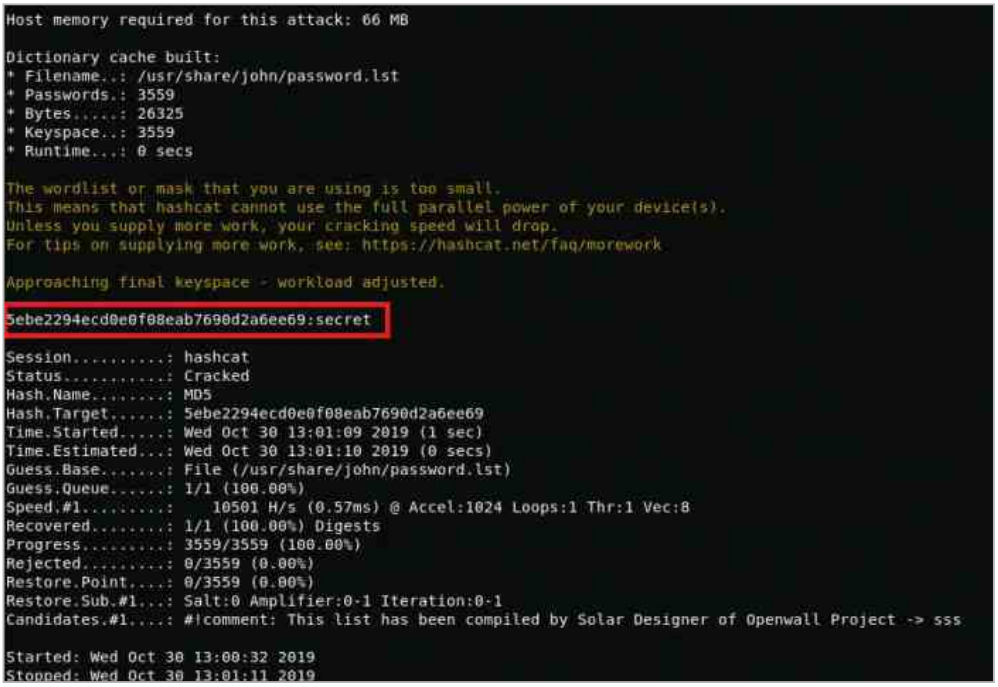

FIGURE 21 *A successful Hashcat operation*

HASHCAT AND SEVERAL OTHER COMMONLY USED CRACKING TOOLS USE A COMBINATION OF TECHNIQUES. HERE ARE A FEW:

Dictionary attacks do exactly what you'd expect, running an abundance of commonly used words to identify any hashes containing weak or common passwords. This is frequently achieved by using the wordlists we mentioned earlier; they're freely available on the Internet and rockyou.txt wordlist is one of the most popular, containing 14 million commonly used passwords.

Hashcat combinator enables two dictionary words to be combined, to create a new list of every word combined with every other word. This detects the use of combination word passwords and also includes the use of special characters to distinguish or separate the words.

A Hashcat mask attack empowers Hashcat to detect the use of certain formats of passwords, such as uppercase letters at the start of a word and a number at the end (very common). This is often dramatically quicker than brute-force attacks, factoring in common human tendencies.

Hashcat rule-based attacks assist users if they have a prior understanding of how the target constructs their password. Then attackers can specify exactly what passwords to try, and can modify, cut, or extend words. For example, if a user knows that a certain special character is used as the second character in an organization's password, a rule-based attack may assist.

Brute-force attacks are often the last resort, permitting a trial-and-error–based approach to cracking passwords, running every combination you can think of. This isn't effective; success is much less likely than with the methods detailed above.

GREAT EXPLOITATIONS: SEARCHLIGHT'S CREDENTIAL BOUNTY AND WHAT IT MEANS

To triage the risk linked to passwords in use today, the Photon team collated and analyzed credentials found in thousands of discrete data breaches from 2016 through 2021, using our proprietary software, SearchLight™. Of course, identifying usernames and (particularly) passwords reliably in very large, poorly formatted files is a considerable challenge. We honed this technique into a fine art, using some simple steps.

First, we identified email addresses. If they weren't present in the file, we gave up on that file; if they were, we looked for passwords. For passwords, the first step was to derive the structure of the file. Often files change structure in the middle of the file, or have badly formatted rows, but SearchLight can process these files into a legible format.

We extracted the username and password from each row, processing only the rows with a „valid" password. We discarded entries like "NULL", "none", etc, and categorized passwords as hashes or plaintext. Then we checked hashed passwords for any format matching our list of known types. (But unknown types might still be included if the field looked like a hash.) Plaintext passwords were flagged, too. For our five-year study period, this process churned out an impressive number of credential pairs.

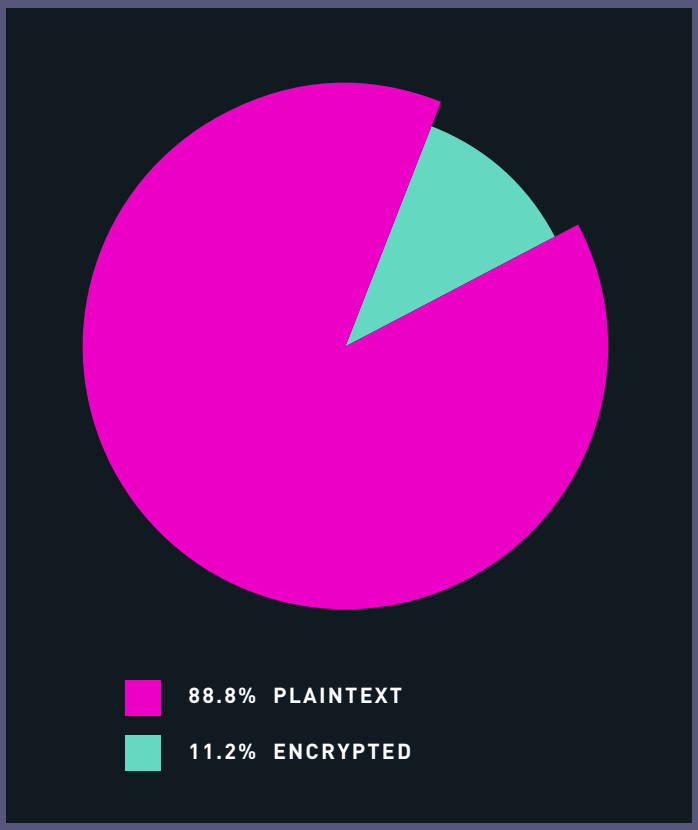

FIGURE 22 *Proportion of plaintext passwords to encrypted*

Password complexity analysis

Ready for that number? The usernames and their passwords identified through SearchLight total 24,649,096,027—that's more than 24 billion credential pairs breached within six years. This staggering number really highlights the scale of the problem. In our 2020 report, we cited 15 billion (up 300 percent from 2018). This new estimate is 64 percent higher than in 2020, highlighting the alarming leap in the number of credentials being exposed; although the overall rate has slowed down, a substantial number of credentials are still being breached every day.

Plaintext passwords make up 88.75 percent of the passwords we identified in our credentials database (see Figure 21). But hold your cringe: Remember that a large number of these passwords were collated from cybercriminal forums that may have used the credential-stuffing or password-cracking techniques we've described. In that case, they would have consolidated them into working username-password pairs before attempting a sale to a third party. So the percentage of hashed passwords is likely to be much higher when they're originally stolen.

BREACHED CREDENTIALS IDENTIFIED	YEAR OF COUNT	INCREASE FROM PREVIOUS REPORT
5,000,000,000	2018	N/A
15,000,000,000	2020	300%
24,649,096,027	2022	64%

Table 1 *Total number of breached credentials identified by Digital Shadows*

POPULARITY RANKING	PASSWORD	NUMBER OF TIMES FOUND
1	123456	30,679,190
2	123456789	17,087,782
3	qwerty	10,589,340
4	12345	10,368,618
5	password	8,987,753
6	qwerty123	5,722,547
7	1q2w3e	5,306,421
8	12345678	5,207,749
9	DEFAULT	4,507,715
10	111111	3,766,387

Table 2 *Top 10 most commonly observed passwords in data set*

THE CURIOUS CASE OF EQ7FIPT7I/Q=

A notable standout in the top 50 most commonly used passwords was not one you could easily guess: EQ7fIpT7i/Q=. In 2013, attackers compromised a backup server of Adobe that was ready to be decommissioned. Adobe had used a single block cipher throughout the breached database, resulting in a weak default password with the same ciphertext in the database: EQ7fIpT7i/Q=. This case shows how different users can have the same weak password created from a complex cipher, which can easily be found out and lead to a breach. The incident reportedly affected 1,911,938 credentials, which was reflected in our data set.

As you'd imagine, sorting through 24 billion-plus credentials is not an unchallenging task. We started by flagging any unique credentials (i.e. unique username-password pairs), after removing duplicates from the total 24 billion credentials. If an account appeared in multiple data breaches, it would only be counted once. We were left with 6.7 billion unique credentials: about 1.7 billion more, or 34 percent more, than in our 2020 research.

The best method to gain meaningful insights from such an enormous number is to whittle it down; we settled on examining only the top 50 most commonly used credentials. The top 50 most common represented a total percentage of 2.40 percent of our 6.7 billion unique credentials. The top 50 is a mix of what you'd expect: almost all are incredibly weak, easily guessable, and related to something the user could easily remember.

We saw references to—or overt use of—the word "password". We saw keyboard-friendly combinations like "qwerty" or "1q2w3e". We saw strings of easily remembered numbers, like 123456…and it's painful to admit that was the most common password. That password actually represented 0.46 percent of our total number of the 6.7 billion unique credentials.

The top ten passwords we found the most are shown in Table 2. Although probably a big portion of these were used for mundane accounts, like a TV or smart thermostat, they're also likely to be in wide use across more-sensitive accounts.

Gauging the strength of brute force

Remember those methods used by Hashcat to conduct offline password cracking attacks? For our research, we used the zxcvbn password strength estimator to determine the average time it takes to successfully crack the passwords of our top 50, in addition to the number of guesses for a brute force attack. Here are the four types of methods we took into consideration, and their definitions[3].

- **Online throttling 100 per hour: An online attack on a service that rate-limits password authorization attempts**

[3] Suman Nepal, Isaac Kontomah, Ini Oguntola, Daniel Wang, "Adversarial Password Cracking", MIT Computer Science and Artificial Intelligence Laboratory, 2019, https://courses.csail.mit.edu/6.857/2019/project/9-Nepal-Kontomah-Oguntola-Wang.pdf

- **Online no throttling 10 per second: An online attack on a service that does not rate limit**
- **Offline slow hashing 10 per second: An offline attack that assumes multiple attackers with user-unique salting and a slow hash function; uses a moderate work factor, such as bcrypt, scrypt, PBKDF**
- **Offline fast hashing 10 (to 1 trillion) per second: An offline attack that uses user-unique salting and a fast hash function, like SHA-1, SHA-256 or MD5; number of guesses per second ranges from one billion to one trillion**

The time needed to crack our top 50 passwords through online throttling is shown in Figure 22, along with the percentage of the group that each segment represents.

Offline cracking methods are—as you would expect—much quicker than online; of course, this varies greatly, depending on the computational strength of the computer conducting the attack, and number of attempts being made per second. From our top 50 passwords, all bar out EQ7fIpT7i/Q= example above were crack-able within a fraction of a single second. Rate-limiting used with an online service also means more time to crack a hash, and ultimately, less chance for a threat actor to compromise an account. Most modern online services use such a process; for those that don't, it's definitely worth starting now.

FIGURE 23 *Time needed to brute-force top 50 passwords via online throttling (100/hr), and percentage of passwords requiring that time*

HOW DOES RATE LIMITING WORK?

Rate limiting runs within an application, rather than on a web server itself. Typically, rate limiting is based on tracking the IP addresses that requests are coming from, and how much time elapses between each request. The IP address is the main way an application identifies who or what is making the request.

A rate-limiting solution measures the amount of time between each request from each IP address, and also the number of requests within a specified timeframe. If there are too many requests from a single IP address within the given timeframe, the rate-limiting solution will not fulfill the IP address's requests for a certain amount of time.

Much in the same way a parent insists their child slow down when eating dinner, a rate limiter instructs unique users to slow down their requests so that the activity is safe. Any user not fulfilling these requests—probably because they're attempting credential stuffing or online cracking—will be blocked.

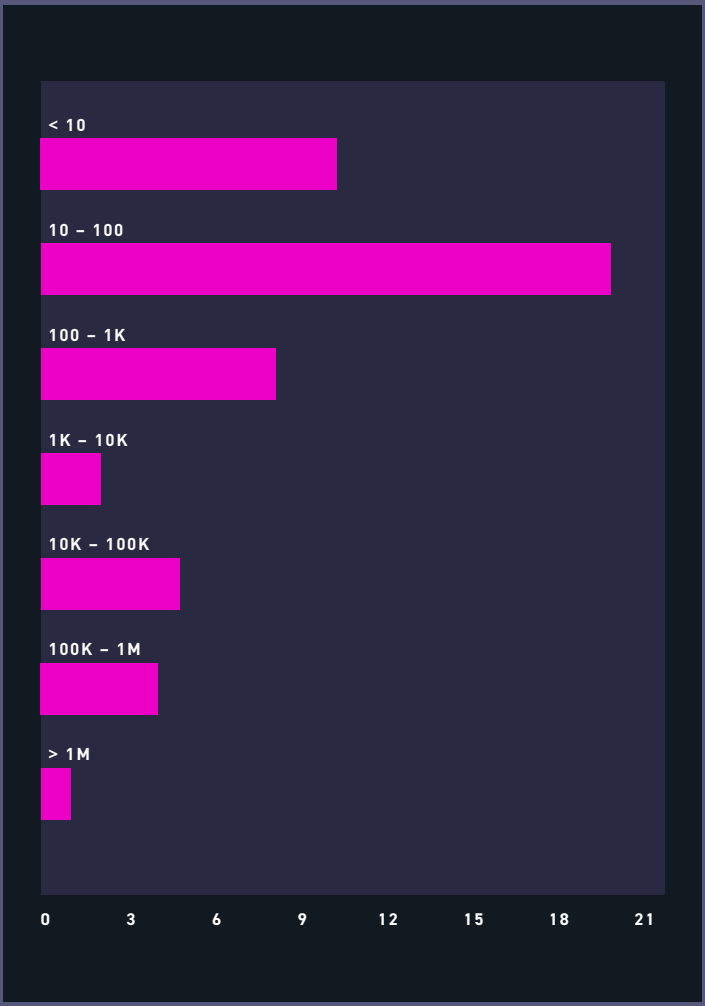

FIGURE 24 *Number of brute-force attempts needed to crack top 50 passwords*

The number of attempts required to bruteforce accounts can be seen in Figure 23. These have been grouped—for example, by those needing fewer than 10 attempts, those needing 10 to 100 attempts, etc. As you can see, most of these top 50 most common passwords can all be cracked with relatively few attempts: 11 percent were cracked in less than 10 attempts, 20 percent were cracked in 10 to 100 attempts, and 8 percent were cracked in 100 to 1,000 attempts. The most needed was about 200,000 attempts, for a six-digit password comprising two digits.

The time it takes to conduct brute-force attempts also varies according to the computational strength of the computer; but, bear in mind that offline attacks run thousands of attempts per second—realistically, all of these passwords would be cracked extremely quickly. If you have a weak password consisting of a common word and a couple of numbers, be aware that even if it's hashed, a skilled attacker will easily be able to crack it.

Still not convinced about the value of sufficient password length and special characters? Take a look at a few password examples we included in Table 3, showing slight differences to track how the changes affect time to brute-force.

Our analysis using zxcvbn showed that although the offline fast hashing method can hypothetically crack a password in less than a second, introducing special characters can draw the cracking time out to several days. One special character (underscore) increased the offline slow hash of London1984 from 3 seconds to 1 hour, 29 minutes, and 21 seconds. Adding two special characters pushed it to 2 days and almost 4 hours. The introduction of several special characters (not shown) essentially made online cracking attempts redundant: Several days would be needed to conduct an online password cracking attack, with or without a limiter.

PASSWORD	NUMBER OF BRUTE-FORCE ATTEMPTS	OFFLINE FAST HASH	OFFLINE SLOW HASH	ONLINE NO THROTTLING	ONLINE THROTTLING
London1984	36,800	0:00:00	0:00:03	1:01:20	15 days, 8:00:00
London_1984	53,610,000	0:00:00	1:29:21	62 days, 1:10:00	22,337 days, 12:00:00
@London_1984	1,868,800,000	0:00:00	2 days, 3:54:40	2,162 days, 23:06:40	778,666 days, 16:00:00

Table 3 *Comparison of time and attempts needed for successful cracking*

STAGE 4: EXPLOITATION

The intrepid ATO attacker has managed to steal your credentials and verify their validity. Their next move is, as you'd imagine, largely limited only to their imagination and initiative. One of the most common means of exploiting stolen credentials is to simply **advertise them for sale,** either after a run through a credential-stuffing tool or password cracker, or in their raw format. Of course, this brings us back to Stage 2 of the attack lifecycle, with one threat actor at the end of the ATO lifecycle facilitating the start for another.

Credential sales typically take place on a dedicated cybercriminal forum or marketplace, but there are other other channels, like the AVCs we detailed earlier. Forum users can make any request, as long as it falls under the rules dictated by the moderators. We've previously reported on what it takes to get banned from such a forum... think "scammy" activity or ignoring the rules. An example of an open request for Spanish and Italian Hotmail user credentials—yep, some people apparently still use Hotmail—can be seen in Figure 24.

The Photon team collates intelligence for our clients from 1.1 billion sources, including 170 million deep and dark web[4] sources—many of them forums/marketplaces. Of course, not all sources are created equal and there are certain pockets of the Internet that prove more fruitful than others.

Three of the cybercriminal forums we use to collate are the Russian forums XSS and Exploit, and the English-language RaidForums. The latter was actually taken down by law enforcement officials in April 2022, but the void will, no doubt, be filled by one or several other forums dedicated to sharing account credentials. The sources used to collate our account database can be seen in Figure 25, which is dominated by those three amigos.

But maybe the holder of stolen credentials doesn't care to sell them. In that case, there's all sorts of other malicious options. ATO is a gateway to sophisticated social-engineering attacks. One type is **business email compromise (BEC):** an escalating threat that has brought significant gains to financially motivated cybercriminals; a recent report from the FBI indicated that the total global financial damage inflicted by BEC activity from 2016 to 2021 equated to $43 billion.

BEC is a fraudulent money-transfer scam targeting companies that conduct wire transfers. This often works by spoofing or compromising corporate and/or publicly available email accounts

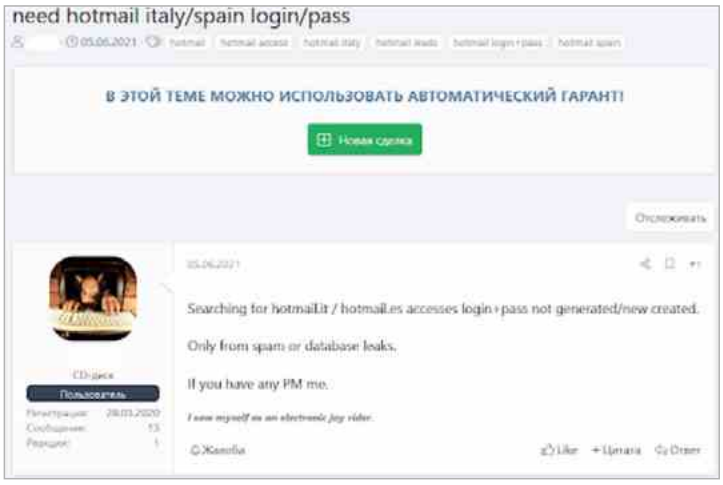

FIGURE 25 *XSS forum user requests Italian or Spanish Hotmail credentials*

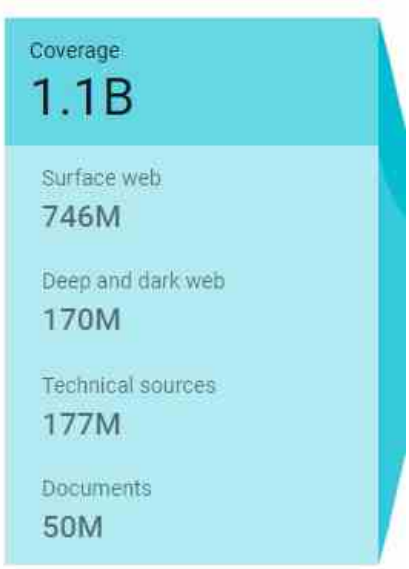

FIGURE 26 *Digital Shadows coverage of 1.1 billion sources*

[4] The deep web refers to a portion of the Internet not indexed by search engines; the dark web is a portion only accessible via certain software/browsers (e.g. Tor).

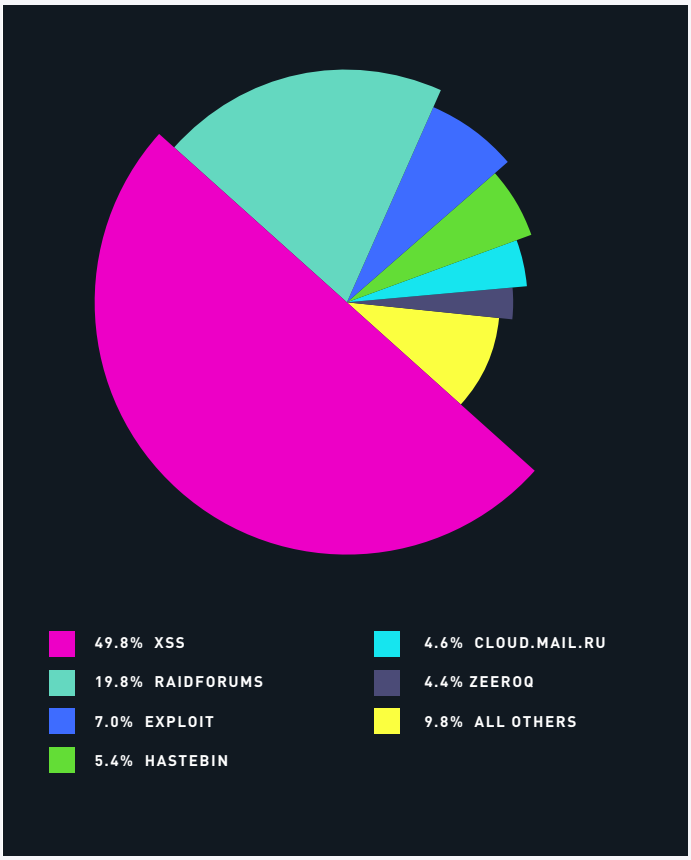

FIGURE 27 *Credential sources that feed the Photon intel database*

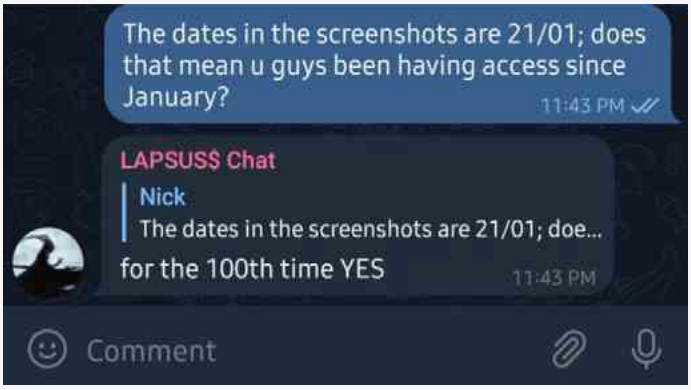

FIGURE 28 *Lapsus$ Group claimed access to Okta beginning 21 Jan 2022 (Source : Futorium)*

of executives, or high-level employees involved in finance or wire transfer payments. Many BEC attacks work by an attacker masquerading as a trusted source and interjecting into existing email threads, but the problem is worsened if an account is compromised.

The risks of ATO have also become very real in recent months since the Lapsus$ Group came on the scene. This collective of data extortionists has fine-tuned the art of socially engineered ATO, compromising high-end technology companies to **steal valuable data** and hold companies to ransom. One ultimate victim was authentication provider Okta, after a compromise of an RDP account used by an Okta contractor. There were concerns that Lapsus$ Group could have accessed data related to thousands of Okta's clients, but Okta clarified that the extortionists' activity had been restricted to 25 minutes and impacted only two customers—despite Lapsus$ Group claiming to have had access since January 2022. Exactly how the RDP account was compromised is unclear, but Lapsus$ Group is known to use many of the social engineering and infostealer malware types detailed in this report, including Redline.

EQUAL-OPPORTUNITY ATO

ATO isn't limited to cybercriminals. Advanced persistent threat (APT) groups, which are often linked to nation-states, frequently abuse credentials to gain initial access to their target companies. Exactly what the long-term access will be used for depends on the groups' motives. Maybe the best recent example of ATO by an advanced threat actor was the supply-chain attack on software giant Solarwinds, which was described as one of the most sophisticated attacks in history. The perpetrators, Russian APT group "NOBELLIUM", used a trojanized software platform update to compromise new networks, but our old friend ATO played a hand in granting initial access—either through social engineering or a brute-force attack.

CREDENTIAL ABUSE MITIGATION: ACCOUNT LOCKDOWN STEPS

Now that you've digested all the potential horrors of ATO, how can you keep accounts safe? Usernames are often easily discovered by a threat actor, and passwords are typically hard for account owners to remember…leading to their picking weak passwords or reusing them among online accounts. Is this a losing battle for security defenders?

Not necessarily, but choose your battles wisely. One way to strengthen passwords and keep them in a solitary, safe location is to **use a password manager**: an application on your phone, tablet or computer that stores your passwords, so you don't need to remember them. Once you've logged in to the password manager using a "master" password, it generates and remembers passwords for all your online accounts. These apps are loaded with technical features to strengthen security, and can even inform you when a current password has been detected as compromised. (Change it immediately!) Choose from several commercially available and free managers, and you'll have taken one of the simplest steps available to reduce the risk of ATO.

Whatever credentials you add to your password manager, make sure they're **sufficiently complex**—let's not forget the Photon analysis showing a correlation between password length and time to crack. For all critical accounts, use passwords longer than ten characters: special characters, numbers, and a combination of uppercase and lowercase letters. Set strong organization password policies to avoid weak or reused passwords, and account lock-out policies to block use if too many failed attempts are conducted.

Multi-factor authentication (MFA), such as two-factor authentication (2FA), has also helped secure the accounts of many online services. An authentication factor is simply a way of confirming your identity, such as via a password, PIN, smartphone, USB key, facial recognition, or fingerprints. For a second factor of authentication, most online accounts favor sending an SMS message or using an authenticator app on a smartphone—even if a threat actor had your username and password, they'd still need your smartphone to access your account. And you'd be alerted to an intrusion attempt when you receive an unexpected authentication request.

Unfortunately, MFA isn't foolproof. Criminals can gain the upper hand with a tactic known as SIM swapping: They take control of a victim's phone number by, essentially, deactivating their SIM and porting the allocated number to an attacker-controlled SIM. They do this by gathering enough information about the victim to successfully social engineer the victim's telecommunications provider to make the swap. You've probably caught wind of this tactic, and its great success, when accounts associated with Jack Dorsey were taken over while he was CEO of Twitter. A host of other celebrities' Twitter accounts have also been compromised.

SIM swapping's main objective is typically to bypass 2FA, so threat actors have advertised numerous methods and services dedicated to that activity. In case you need yet another wake-up call: Don't sleep through this threat. Keep personal details private, including your date of birth, pets' names, family names—anything that can be used to answer security questions that verify identity.

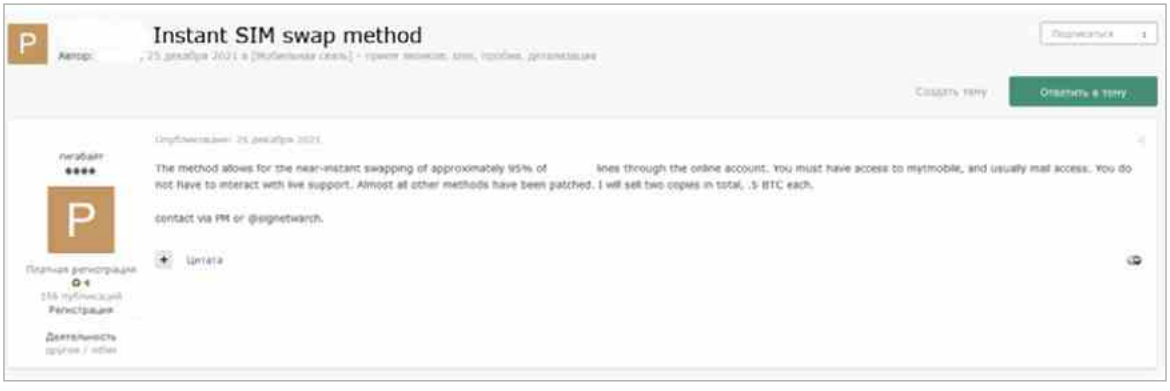

FIGURE 29 *Exploit user detailing SIM swapping methods*

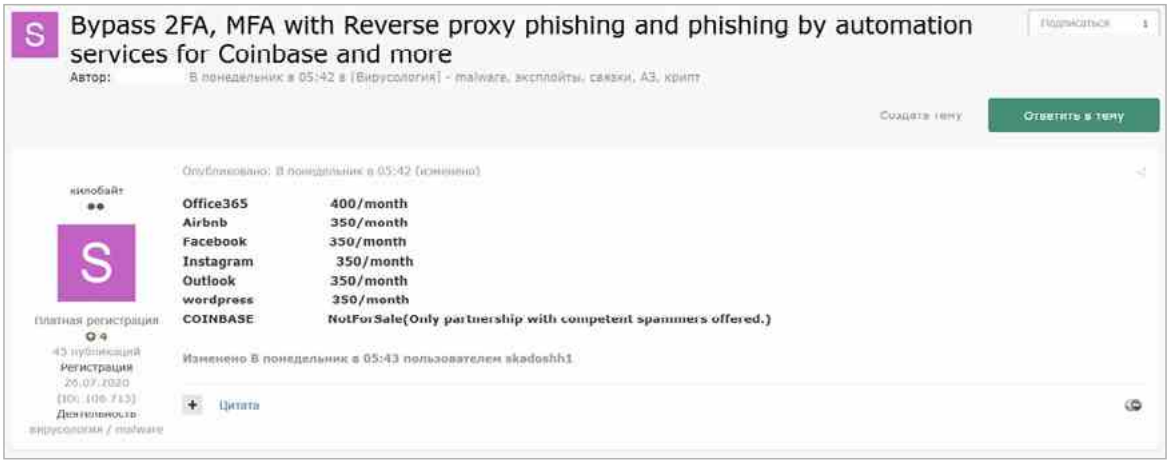

FIGURE 30 *Exploit user advertising service for bypassing 2FA*

To overcome SIM swapping, a great alternative to MFA is an **authenticator app.** It generates a new, random six-digit code every 30 seconds, which a user must enter on the website they're trying to access. Many pages ago, we hinted at a potentially utopian passwordless future; it's true, the world will probably move toward greater reliance on safe, passwordless authentication, such as biometrics. But at this moment in time, no authentication method is 100 percent safe; we must remain on high alert for phishing and social engineering attacks. (To read more about SIM swapping specifically, see a great Photon blog on this topic from March 2022.)

A bot management service can help you defend against credential stuffing and online attacks against your services. That kind of service combines rate limiting with an IP reputation database to stop malicious bots from making login attempts. Bot management services can also implement machine-learning techniques to learn behaviors of threat actors attempting to cover attack methods.

That's just about it, but we can't cover mitigation without mentioning the importance of **monitoring credential compromise**. Keeping track of when an employee's credentials may have been compromised will help network defenders respond quickly. Consistent and accurate monitoring can be difficult to maintain, and may not always prevent credential abuse, but swift action following such abuse can minimize the risk and stop a threat actor in their tracks.

ADMITTING YOU HAVE A PROBLEM IS ALWAYS THE FIRST STEP TO SOLVING IT, AND YOU'RE NOW FULLY AWARE OF THE DANGERS OF ATO. UNTIL WE LIVE IN A WORLD THAT'S BYPASSED PASSWORD USE ENTIRELY, WE'LL BE KEEPING OUR GUARD UP, SECURING OUR PASSWORDS AND SCANNING THE HORIZON FOR THAT INTREPID ATO ATTACKER.

About Digital Shadows

Digital Shadows minimizes digital risk by identifying unwanted exposure and protecting against external threat. Organizations can suffer regulatory fines, loss of intellectual property, and reputational damage when digital risk is left unmanaged. Digital Shadows SearchLight™ helps you minimize these risks by detecting data loss, securing your online brand, and reducing your attack surface.

To learn more and get free access to SearchLight™, visit
www.digitalshadows.com

London, UK **San Francisco, CA** **Dallas, TX**

Authors: Digital Shadows Photon Research Team

digital shadows_