Q2 2019 DataVisor Fraud Index Report

MINING DATA TO EXPOSE THE DIGITAL FOOTPRINTS OF FRAUD.

DATAVISOR

Table of Contents

Foreword

Fraud losses have reached staggering levels, and while there continue to be minor fluctuations year-over-year, the overall situation is dire: in 2018 alone, fraud losses hit $14.7 billion.

Fraud losses have reached staggering levels, and while there continue to be minor fluctuations year-over-year, the overall situation is dire: in 2018 alone, fraud losses hit $14.7 billion. Many different attack types contribute to these numbers, but Account Takeover (ATO) is uniquely devastating, accounting for $4 billion of those 2018 losses. In the e-commerce sector, nearly 40% of all fraud losses in 2018 were due to identity theft and synthetic identities, and this represents almost a 100% increase over the preceding year.

Account compromise come in many forms, with one of the most common being credential stuffing. Given how often data is exposed in breaches, it's not surprising fraudsters are using all that data to try and determine credential validity through brute force attacks. According to Ponemon Institute's The Cost of Credential Stuffing Report, companies experience 12.7 credential stuffing attacks each month, with more than 1,200 user accounts being typically targeted in each credential stuffing attack. Approximately 12.4 percent of these attempts are successful.

As mobile phones become an increasingly common part of our identity (e.g., phone numbers used as logins, phones as the primary factor for texts, voice, or other types of multi-factor authentication), fraudsters have also shifted their focus to more aggressively target mobile accounts. Mobile phone account takeovers rose nearly 180% from 2017 to 2018, resulting in nearly 700k ATO incidents. Hijacking a phone number means that the fraudster not only controls all online accounts tied to the number but can also intercept SMS messages—a preferred method for verifying financial account logins.

As wretched as these numbers sound, they only paint a portion of the picture when it comes to addressing the challenge of ATO attacks. Not only do businesses have to defend against the bad actors, but they also have to simultaneously protect their good customers. If a good customer's account gets hijacked, they need to have confidence that they will be protected before any damage can occur. So it's not enough to rely on a fraud solution that addresses only the end action—the actual theft of assets. Businesses have to focus on the good users as well and address their account issues before the next crime occurs.

This is easier said than done. It requires proactive detection. You have to spot potential attacks and stop them before they can launch. You have to be able to identify incubating accounts, recognize what they're being primed for, and neutralize them before they can be harnessed for use in a major coordinated attack. Too many existing solutions address fraud at the transaction level. However, with ATO, that's already too late. Successful ATO prevention necessitates prevention at the account level—you need to know the moment an account gets compromised, so you can prevent damage, and preserve the user's safe and secure experience.

- Ting-Fang, Director of Research, DataVisor

THE DATAVISOR GLOBAL INTELLIGENCE NETWORK

The DataVisor Global Intelligence Network (GIN) leverages deep learning technologies to provide real-time, comprehensive digital intelligence based on a vast set of data signals that include IP addresses, geographic locations, email domains, mobile device types, operating systems, browser agents, phone prefixes, and more. All told, the GIN aggregates anonymized signals across a global client database of more than four billion users.

By analyzing the connections between these data points in context—not just in isolation—DataVisor provides fine-grained signals and reputation scores that can be consumed directly in detection, or used to enhance rules engines and machine learning solutions.

To produce this report, we processed and analyzed the following for the period January-March, 2019:

52 billion events

1.1 billion users

427 million IP addresses

4.77 million /24 IP subnets

1.46 million email domains

5.84 million user-agent strings

263K device types

481K phone number prefixes

SECTION 1

How Are Accounts Compromised?

There are seemingly an almost unlimited number of ways a fraudster can compromise an online account, but certain techniques are particularly prevalent. They run the gamut from phishing and brute-force attacks, to banking trojans and mobile phone hijacking.

PASSWORD SPRAYING

These attacks are a type of "brute-force" attack, meaning they require no particular degree of sophistication, and are essentially accomplished through trial-and-error at a large scale. In this type of brute force attack, fraudsters—usually relying on scripted bots—spray relentless pairs of common usernames and passwords in hopes of landing on the right combination to enter an account. Users with weak passwords and generic usernames are particularly vulnerable. Once a fraudster gets access to an account, they will use it as long as they can to commit fraudulent acts and to drain the compromised account of all value.

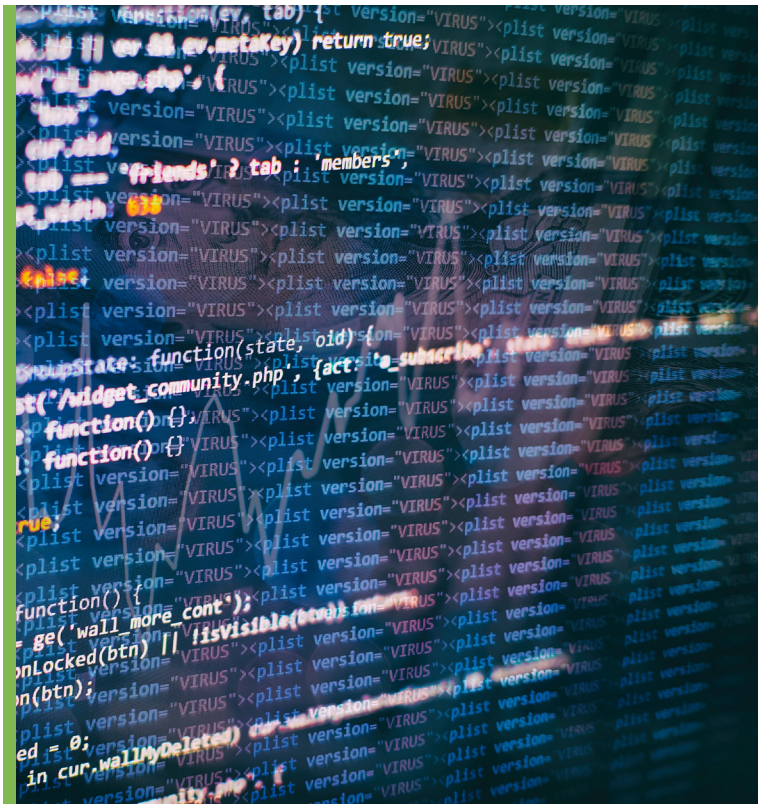

CREDENTIAL STUFFING

As discussed previously, credit stuffing attacks are becoming increasingly common. In a credential stuffing attack, fraudsters leverage massive troves of leaked legitimate user credential data to begin firing pairs of names and passwords at other sites in hopes of getting a "hit"—an instance in which a particular combination works. Users that reuse passwords across sites are particularly vulnerable to these types of attacks.

According to a recent State of the Internet Report from Akamai, there were nearly 30 billion credential stuffing attacks in 2018, with hundreds of millions of attempts taking place every day.

Credential stuffing attacks often rely on readily available tools to automate the process. SNIPR, for example, is a popular entry-level tool that includes predefined configurations for popular websites, includes proxy support and community forums.

SOCIAL ENGINEERING

As prevention techniques adapt to ongoing fraud, it follows that new attack types will emerge while others become obsolete. Certain things, however, never seem to change. Phishing, for example, remains as omnipresent as ever. According to the latest Verizon Data Breach Investigation Report (DBIR), 32% of data breaches involved phishing as part of the attack tactic. As noted previously, mobile is particularly vulnerable to fraud, and this holds especially true with regards to phishing attacks. As reported by DarkReading, small screens and limited security measures are likely causes of heavy phishing activity on mobile channels, and nearly 60% of all mobile fraud attacks are mobile phishing.

MALICIOUS SOFTWARE

Keyloggers, trojan viruses, spyware, and various other types of malicious software are used by fraudsters to intercept or harvest sensitive information. Banking trojans are particularly dangerous, as they're used to steal financial credentials and drain bank accounts. Many banking trojans work by overlaying a "fake" login page on top of a legitimate bank website. When a bank customer logs in, believing the page to be authentic, the credentials are intercepted and stolen. This user information is sent back to the fraudster behind the trojan, who then uses the credentials in criminal attacks. All of this takes place without a user realizing what's transpired—by the time the evidence is obvious, the crimes have already been committed, and the money already stolen. Banking trojans spread primarily through spam or phishing emails. A recent study from ProofPoint showed that banking trojans are found in 56% of all malicious emails, with the Emotet malware making up 76% of all banking trojans.

PHONE HIJACKING

The widespread adoption of SMS messages for second-factor authentication has not stopped fraudsters from taking over accounts.

In a SIM swapping attack (also called SIM hijacking), a fraudster asks a mobile carrier to switch a phone number to another SIM card under their control by impersonating the actual account owner. This type of account takeover gives the fraudster access to all online accounts tied to the phone number as well as incoming SMS messages, allowing them to easily bypass second-factor verification measures often used to protect sensitive accounts.

Other methods exist for intercepting text messages, including posing as rogue public wi-fi hotspots or fake cell towers. A more elaborate scheme exploits vulnerabilities in the SS7 routing protocol, which is used by mobile networks to route calls and texts. A vulnerability in the protocol allows anyone with access to a gateway on the SS7 network to intercept calls and texts or to track specific devices, even from a remote location.

SECTION 2

Financially Motivated ATOs

ATO attacks are complicated to execute and difficult to detect. This is because they're multi-step processes.

The criminal behind the attack must first obtain valid user credentials—either by direct theft or through brute force, trial-and-error efforts like credential stuffing. Once in, they can take several actions. They can drain the account of any value it might contain. They can use the account to register for additional services or benefits, then leverage those registrations for further criminal activity. They can get bank loans, open new accounts, post fraudulent listings, and more. Here are some of the most common downstream attacks that are the result of compromised accounts:

Financial fraud: Examples include unauthorized withdrawals or fraudulent transactions using on-file credit and debit cards.

Spam: Goals include spreading scams, fake news, or malicious links. Spam can appear anywhere that accepts user-generated content, including discussion forums, direct messages, and reviews and ratings sites.

Phishing: Attackers can assume a compromised user's identity and launch phishing attacks on others in their community to steal their credentials, personal information, or sensitive data.

Promotion abuse: Fraudsters can take advantage of promotions available on e-commerce sites to purchase discounted items in bulk—preferably those that are easily transferable—such as virtual currency, rewards points, prepaid cards, and more.

Card testing: Fraudsters can make small purchases, or attempt to add credit cards to compromised user accounts, to check the validity of stolen credit cards. Once a card is determined to be valid, it can be used for further criminal activity.

Virtual currency fraud: Virtual currencies that are worth real money include reward points, promotional credits, and in-game virtual items, all of which can be harvested for real-world gains.

SECTION 3

The Anatomy of Account Takeover

For this report, we analyzed over 50K compromised accounts and approximately 100 detected fraud campaigns across multiple global online services.

MAIN FINDINGS

Our main findings include:

- **Most account takeover attacks go unnoticed.** The majority of compromised accounts are dormant accounts into which the user has not logged in for an extended period of time. 65% of compromised accounts have not logged in for more than 90 days.
- **Fraudsters that compromise financial accounts take additional steps to stay under the radar.** 20% of compromised accounts were accessed within 300 miles of the account owner's location. This makes fraudulent activities less likely to trigger suspicion since they do not deviate significantly from the account owner's normal activities.
- **After compromise, fraudsters move quickly.** 72% of financial accounts made fraudulent transactions within one hour of compromise.
- **ATO attacks are conducted at scale.** The majority of successful ATOs come from password spraying or credential stuffing attacks where hundreds of thousands of unique IP addresses are used for logging in to user accounts via bots and automated scripts.

DORMANT ACCOUNTS

The large majority of compromised accounts are in a dormant state at the time they are fraudulently accessed, meaning they are owned by users who have not logged in for an extended period of time. 65% of these accounts belong to users that have not logged in for more than 90 days, and 80% of these accounts belong to users that have not logged in for more than 30 days. ATOs involving dormant accounts are difficult to detect. The takeover (and often the subsequent fraudulent activity) usually goes unnoticed by the dormant user, as they are not actively managing their account. Additionally, the online service where the account is registered may not have enough information about the user to detect that there is a change in the account behavior. Without a track record of activity, it is more challenging to identify suspicious anomalies.

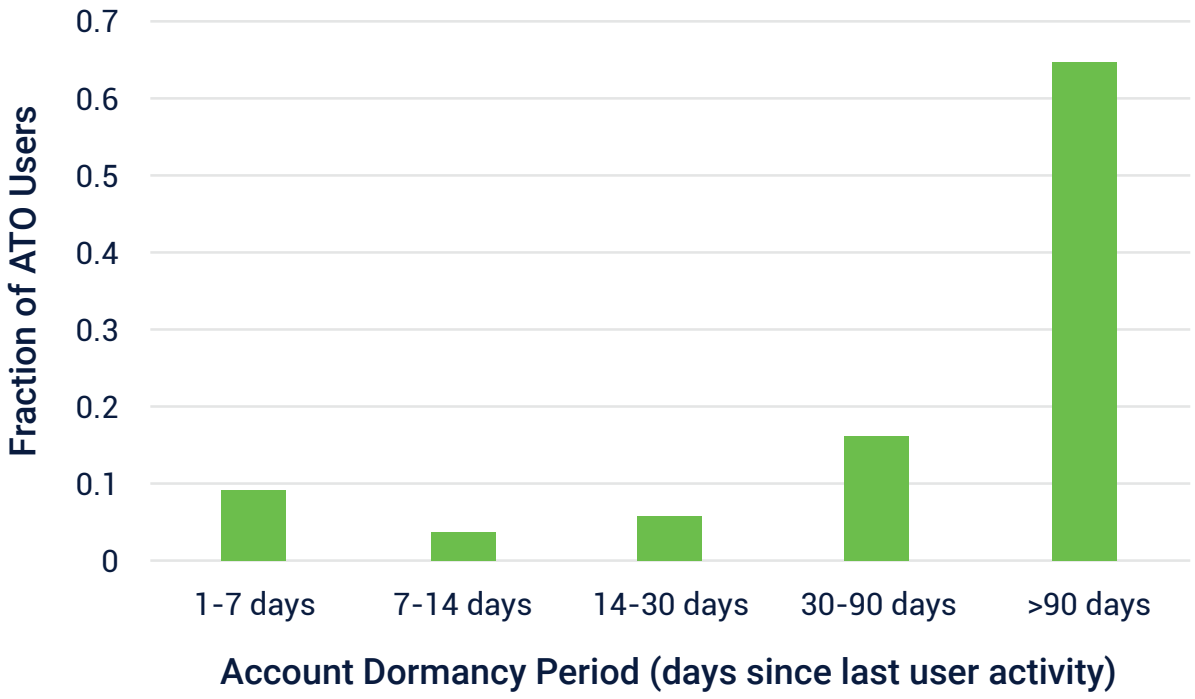

Figure: The number of days between last user activity and account takeover. Most compromised accounts are dormant accounts that belong to users who have not logged in for an extended period of time.

STEALTHY BEHAVIORS

Once an account has been compromised, the odds of detecting it increase if the fraudster's actions deviate significantly from the account owner's previous behavior.

A typical example of this kind of detection strategy is when a purchase is flagged as suspicious if it is made in a location where the user has never made a purchase before, particularly if it's notably far from the geographic area where the user is normally active.

To avoid this kind of detection, fraudsters—especially those controlling compromised financial accounts—will attempt to login "close to" the account owner's location. This can make their activities appear more legitimate. In our data, 20% of compromised financial accounts were accessed within 300 miles of the account owner's known location, while this is true for only 3.4% of compromised accounts on social platforms.

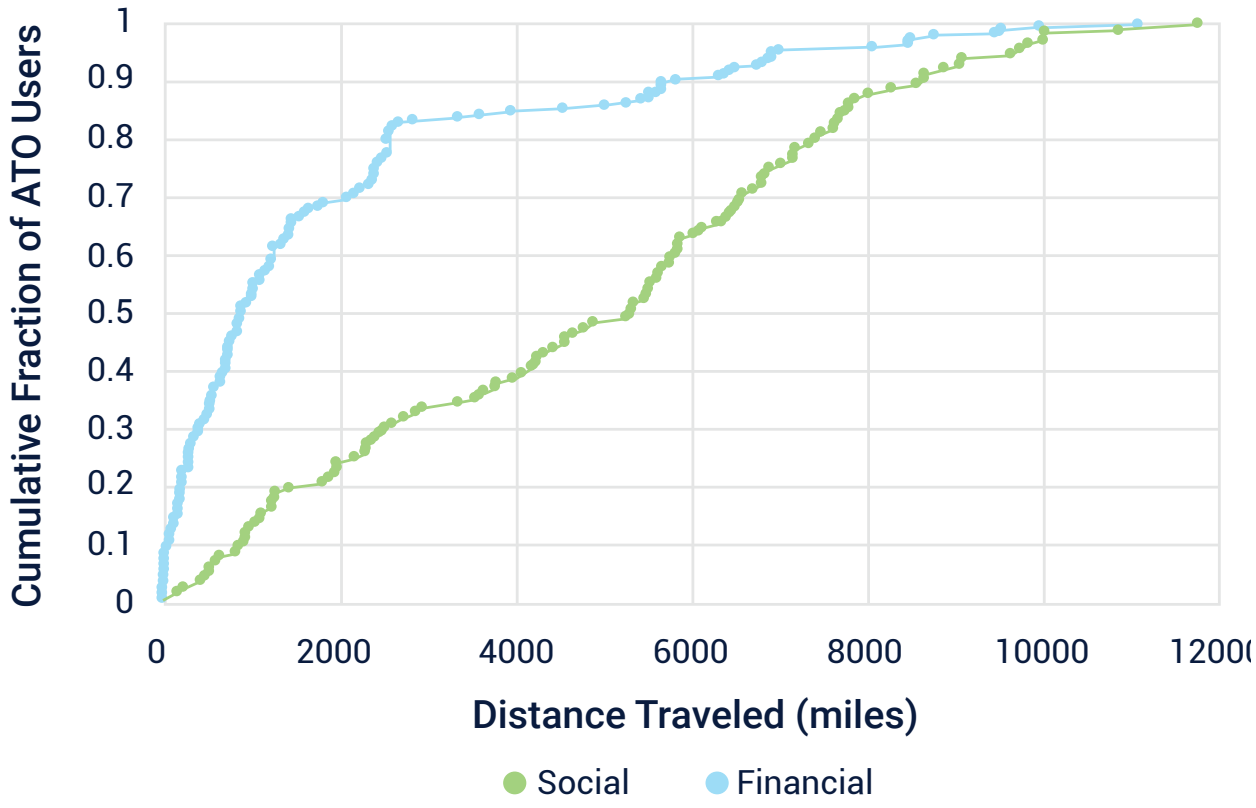

Figure: The distribution of the distance (in miles) between the last user activity and the ATO login for compromised accounts on financial and social platforms. Compromised financial accounts tend to stay "closer" to the account owner's original location.

Fraudsters often try to act quickly after obtaining access to an account. We found that 72% of compromised financial accounts are used to execute fraudulent transactions within one hour of the initial login. Financial attacks are under higher time-constraints since stolen financial information (credit card numbers, banking information, and more) expires quickly.

By contrast, compromised accounts on social platforms tend to be less time-sensitive. In one large ATO attack on a social platform, 81% of the compromised accounts did not start posting spam or scam messages until three weeks after the initial login. This kind of strategy—known as "account incubation"—makes detection especially hard. Reactive detection can spot the fraudulent activity, but by that time, the damage is already done. Proactive strategies, on the other hand, can surface patterns indicating the presence of incubating accounts that are being primed for future use in a coordinated and large-scale attack.

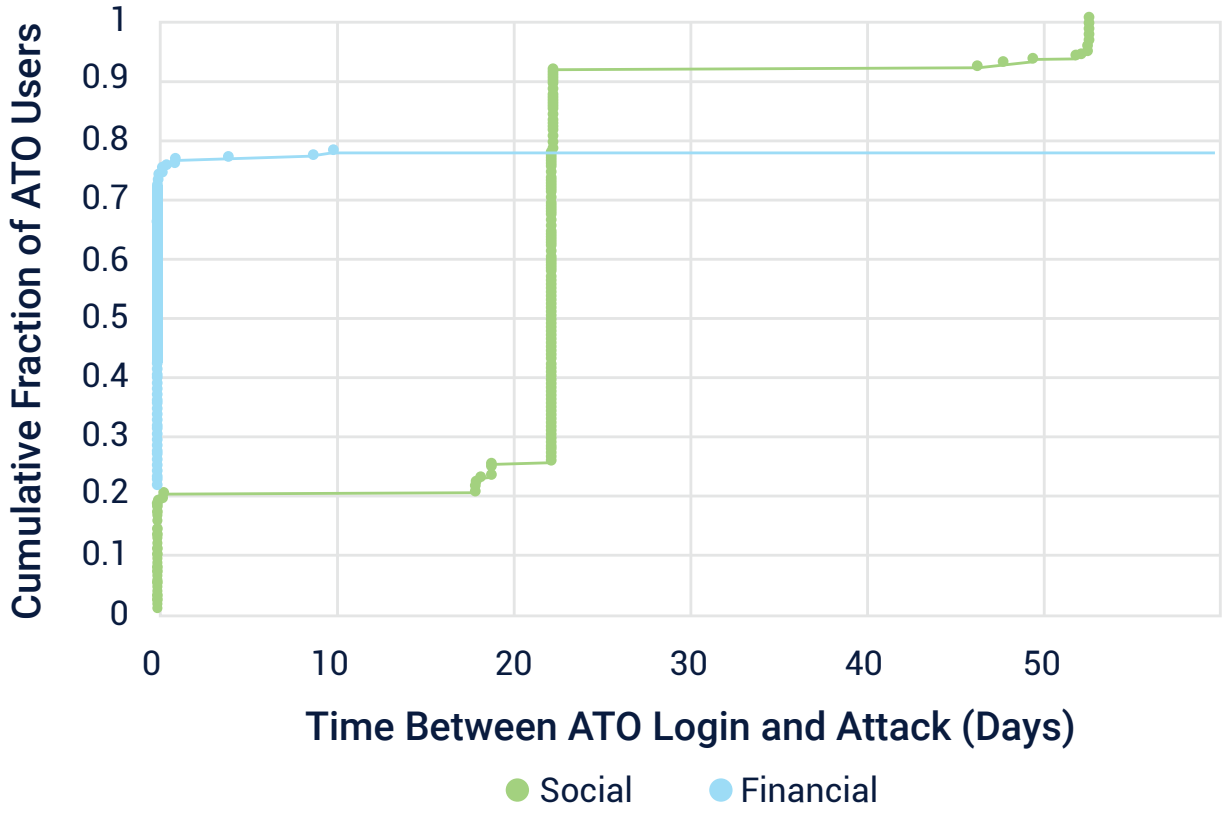

Figure: The distribution of the time (in days) between the initial ATO login and the first attack. Compromised accounts on financial platforms are used much quicker than those from social platforms.

COORDINATED ACCOUNT TAKEOVER

ATO attacks are conducted at scale. The majority of successful ATOs are the result of password spraying or credential stuffing attacks, where hundreds of thousands of unique IP addresses are used for logging in to user accounts via automated scripts.

The maps below show the locations of IP addresses where four different account takeover attacks originated. Each attack involved tens of thousands of compromised accounts. The blue dots indicate the initial ATO login locations, while the green dots indicate the origin of the subsequent attack. The distributed nature of these large-scale, coordinated attacks shows that the fraudsters have a large amount of IP resources at their disposal; likely from botnets consisting of compromised machines.

None of the IP addresses that were used to access compromised accounts were found in publicly available blacklists containing spamming hosts or malware servers **(e.g., https://iplists.firehol.org/)**. This suggests that account takeover attacks may utilize slightly different infrastructure from other types of cybercrime.

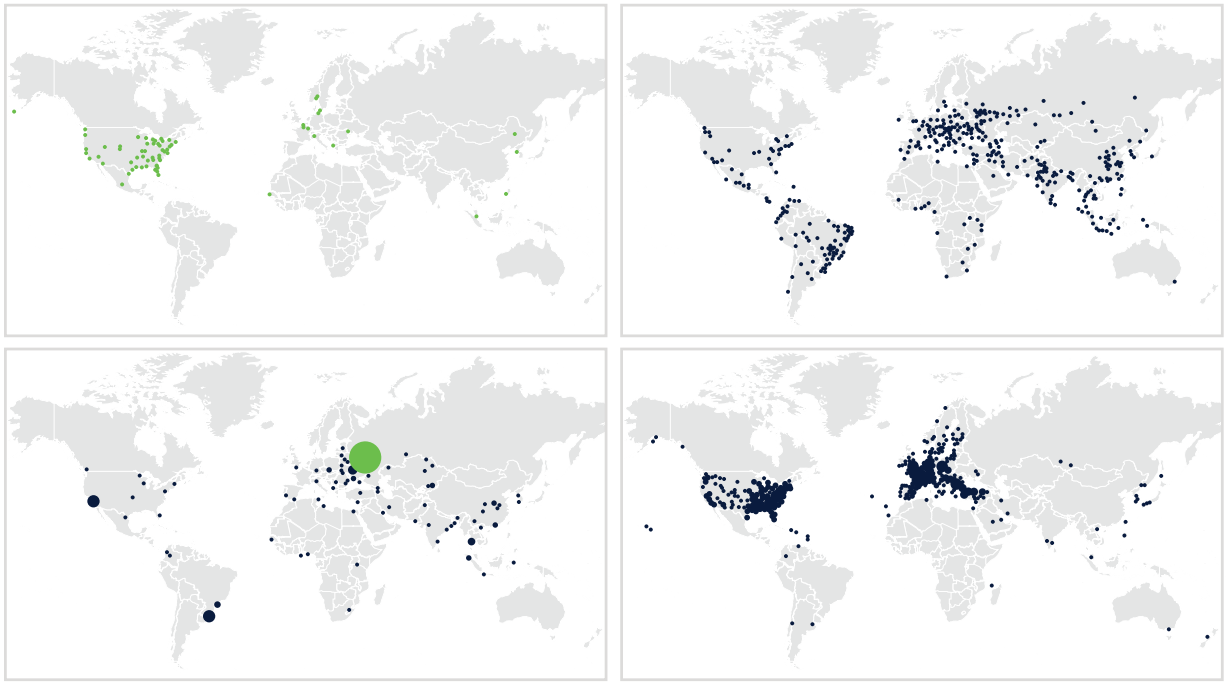

Figure: Locations of IP address from which account takeover attacks originated.

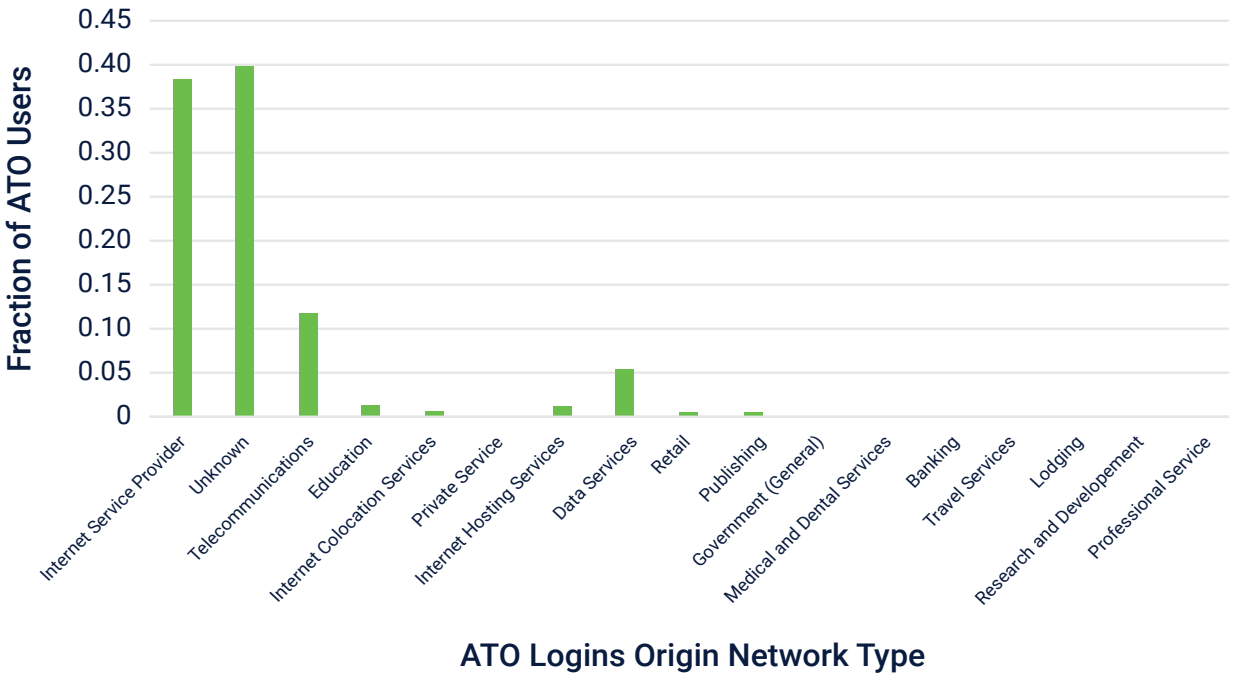

Figure: The type of IP networks from which the ATO activity originated. Over half of the compromised accounts were logged in from networks associated with internet service providers or telecommunication providers. These are likely botnets consisting of compromised machines.

Leveraging distributed IP addresses across the world likely helped the fraudsters evade detection since each IP address was only used to access a handful of compromised accounts. However, in some cases, fraudsters go even further to make their activities appear "normal" or "random."

In one such attack, the fraudster manipulated their connections with the online service so that each fraudulent account used a different user-agent string (a text string that identifies an HTTP client—such as a browser—to the web server, specifying the client's OS and browser version, among other software configurations). Though each user-agent string was unique, all of them followed the same pattern—just the phone model was replaced.

Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**LG-P712**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**SM-P555S**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**GT-S7262**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**HUAWEI-CUN-L01**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**SPH-M920**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**SM-G350M**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**SM-S550TL**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**LG-P970**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**LG-V521**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**HUAWEI-NMO-L31**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**Ideos**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**LG-P505R**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**SM-C5000**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**LG-E440**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**SM-J727T**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**dtab01**)
Dalvik/2.1.0 (Linux; U; Android 5.1.1; Build/**SM-G611F**)

Figure: Examples of manipulated user-agent strings that are "different" but still contain the same pattern. The section of the user-agent string that varies among users is highlighted in bold.

AUTOMATED ACCOUNT TAKEOVER

There appear to be two general types of ATOs. One is fully automated. Using scripts or tools (like SNIPR, https://snipr.gg/), the fraudster launches a massive number of login attempts to access victim accounts. Unlike normal human activities, which exhibit diurnal patterns corresponding to awake/sleeping hours, the scripted nature of the fraudulent activities means that they can take place at all hours of the day, consistently. Approximately 80%-90% of the fraudulent ATO logins fall into this category.

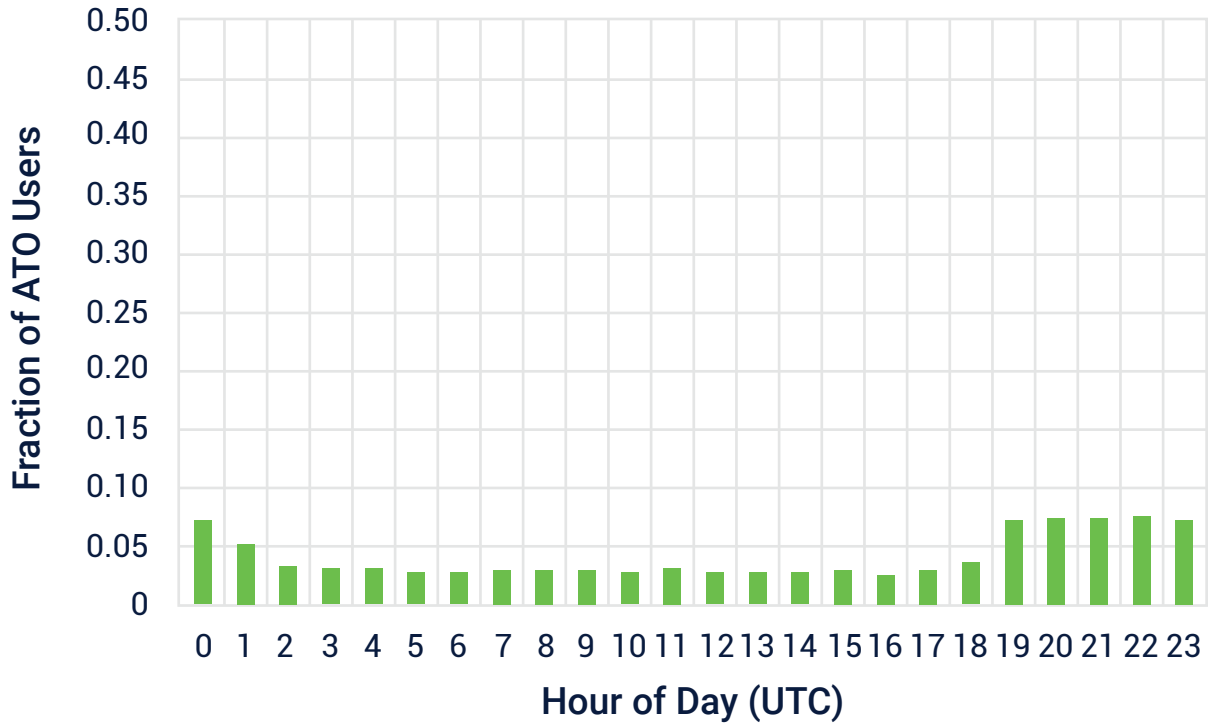

Figure: The distribution of time-of-day when compromised accounts were accessed by the fraudsters. This shows the type of ATO where fraudulent accesses are completely automated.

The other type of ATO attack is only partially automated. In the example shown below, we observe fraudulent accesses concentrated around specific periods of the day (e.g., from 10am until 3pm), with attack activities (e.g., fraudulent transactions or spam) taking place sporadically (e.g., mostly at 5pm, then 9pm). It is likely that fraudsters gain initial access to the accounts through automated scripts, and only perform subsequent attack actions manually after access is confirmed.

This second type of ATO attack reflects cases where the fraudster conducted the attack in stages, or the compromised accounts changed hands. For example, the ATO attack could be conducted by crime rings specialized in obtaining access to accounts (e.g., via credential stuffing). The successfully compromised accounts are then sold off to another party that uses them for various downstream attacks, including spam, phishing, fraudulent transactions, and more.

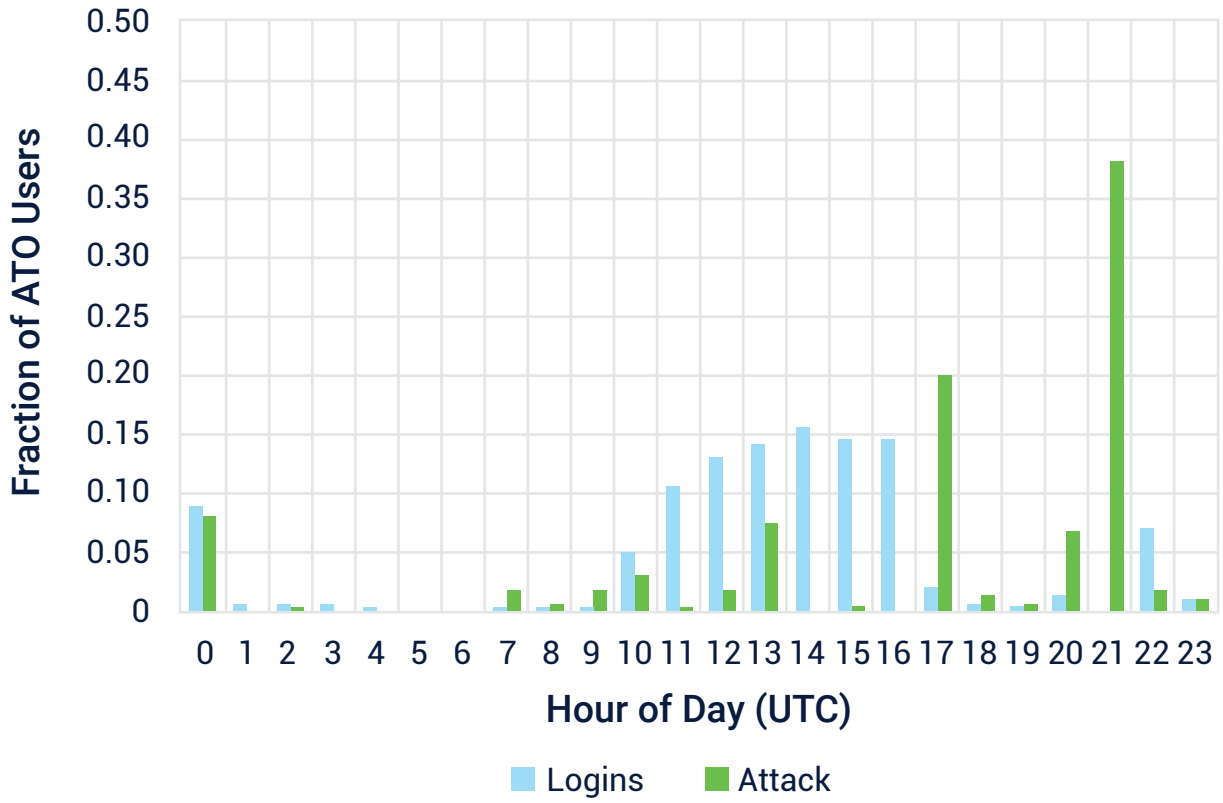

Figure: The distribution of time-of-day when compromised accounts were accessed by the fraudsters. This shows another type of ATO where the attack is only partially automated. While fraudsters gained initial access to the accounts through automated scripts, they performed subsequent attack actions manually after access is confirmed.

SECTION 4

Account Takeover Prevention and Remediation

Dealing with ATO fraud presents a unique challenge in that it requires addressing potential attacks as much as it does actual attacks.

A data breach in and of itself does not cause actual damage. However, as soon as that data leaks, you have to start thinking about what could happen should that data fall into fraudster hands. When it does, you have to try and predict how criminals will use that data. What kind of attack, which targets, and to what ultimate purpose? Even the act of taking over an account isn't a definitive threat; if the fraudster does nothing with the compromised account, or if there is no value to drain from the account, there is no real damage. But the potential for damage is enormous.

The challenge of dealing with ATO fraud is unique in other ways as well. For example, it's not just about dealing with the bad actors—detecting, exposing, and neutralizing them. You also have to support your good users. As a business, you must preserve exemplary experiences for your good customers. Ideally, their accounts are always secure. However, if an account does get hijacked, a customer needs to know their assets will be safe and protected before any damage can happen.

Account takeover is at minimum an extremely frustrating experience to endure, and when it succeeds to the full measure of its potential, the results can be devastating—for businesses and their customers both.

Proactive Fraud Management with DataVisor

Early detection is the difference between damage response and damage prevention. For too long, legacy fraud solutions have been reactive, responding to attacks after the fact, and hoping at best to limit losses. Modern digital fraud has become too sophisticated for this approach to remain viable. Modern fraudsters are agile, adaptive, and technologically adept. They marshall armies of bots to commit fraudulent actions at a massive scale, and their attacks are coordinated, complex, and global. Rampant data breaches provide them with a study diet of ill-gotten personal and financial information, and ongoing platform vulnerabilities across industries offer easy portals to profit.

Fortunately, there is an answer to these challenges. DataVisor is leading the way in delivering production-ready unsupervised machine learning solutions that enable proactive fraud prevention. When combined with big data architecture, global intelligence resources, holistic data analysis and management, and contextual detection capabilities, proprietary algorithms surface correlated patterns of suspicious activity early, before attacks are launched and damage occurs. Sophisticated fraud models reveal connections between accounts and actions that go undetected when viewed in isolation, and malicious accounts are neutralized before they can be brought out of incubation and used in an attack.

Comprehensive fraud management with dCube

dCube is a comprehensive fraud management solution combining transformational AI-powered technology with a streamlined workflow to enable large enterprises to proactively thwart both known and unknown fraud. dCube features a hyper-modern architecture built to manage complex digital signals and behavior analytics using the most advanced machine learning technologies at big data scale, empowering large enterprises to identify and prevent even the most sophisticated attacks. dCube facilitates unparalleled agility by allowing all stakeholders to collaborate on a single platform, eliminating organizational bottlenecks and enabling real-time detection and response.

Managed fraud services with dVector

dVector is a best-in-class managed fraud detection service powered by transformational machine learning technology. It provides optimized detection scores with clear and actionable reasons in real time so organizations can take action against known and unknown fraud before damage occurs. dVector efficiently handles structured and unstructured data at large scale and leverages unsupervised and supervised machine learning technologies to automatically identify and categorize different types of fraud and abuse. dVector delivers fully optimized detection scores alongside clear reasons that detail how each instance of fraud or abuse is committed, empowering businesses to take decisive and proactive action with confidence.

Conclusion

The true end goal of any fraud management strategy isn't actually detection; it's prevention. Achieving this goal requires action from all stakeholders—businesses, individuals, and fraud management solution providers. Companies need to work with fraud management solution providers to adopt future-facing solutions powered by AI and machine learning, that can proactively detect brewing attacks before they launch and can cause damage, and which can provide ongoing monitoring for incubating malicious accounts. Simultaneously, customers need to consistently follow best practices for account security, including enabling multi-factor authorization and leveraging password managers.

The good news is that proactive fraud management is a reality. Unsupervised machine learning algorithms can uncover hidden correlations and reveal suspicious patterns in real time, without the need for historical labels, extensive training periods, or time-consuming retuning. Contextual detection capabilities and holistic data analysis make it possible to address ATO at the account level instead of the transaction level, stopping attacks before damage happens. Scalable detection engines can keep pace with the scope of modern bot-powered attacks, reviewing thousands of fraud signals to surface even the most cleverly-disguised actions. No matter how sophisticated and how technologically advanced fraudsters may be, and no matter how massive and widespread their fraudulent activities, they still leave digital footprints that can be revealed with advanced AI-powered fraud management solutions.

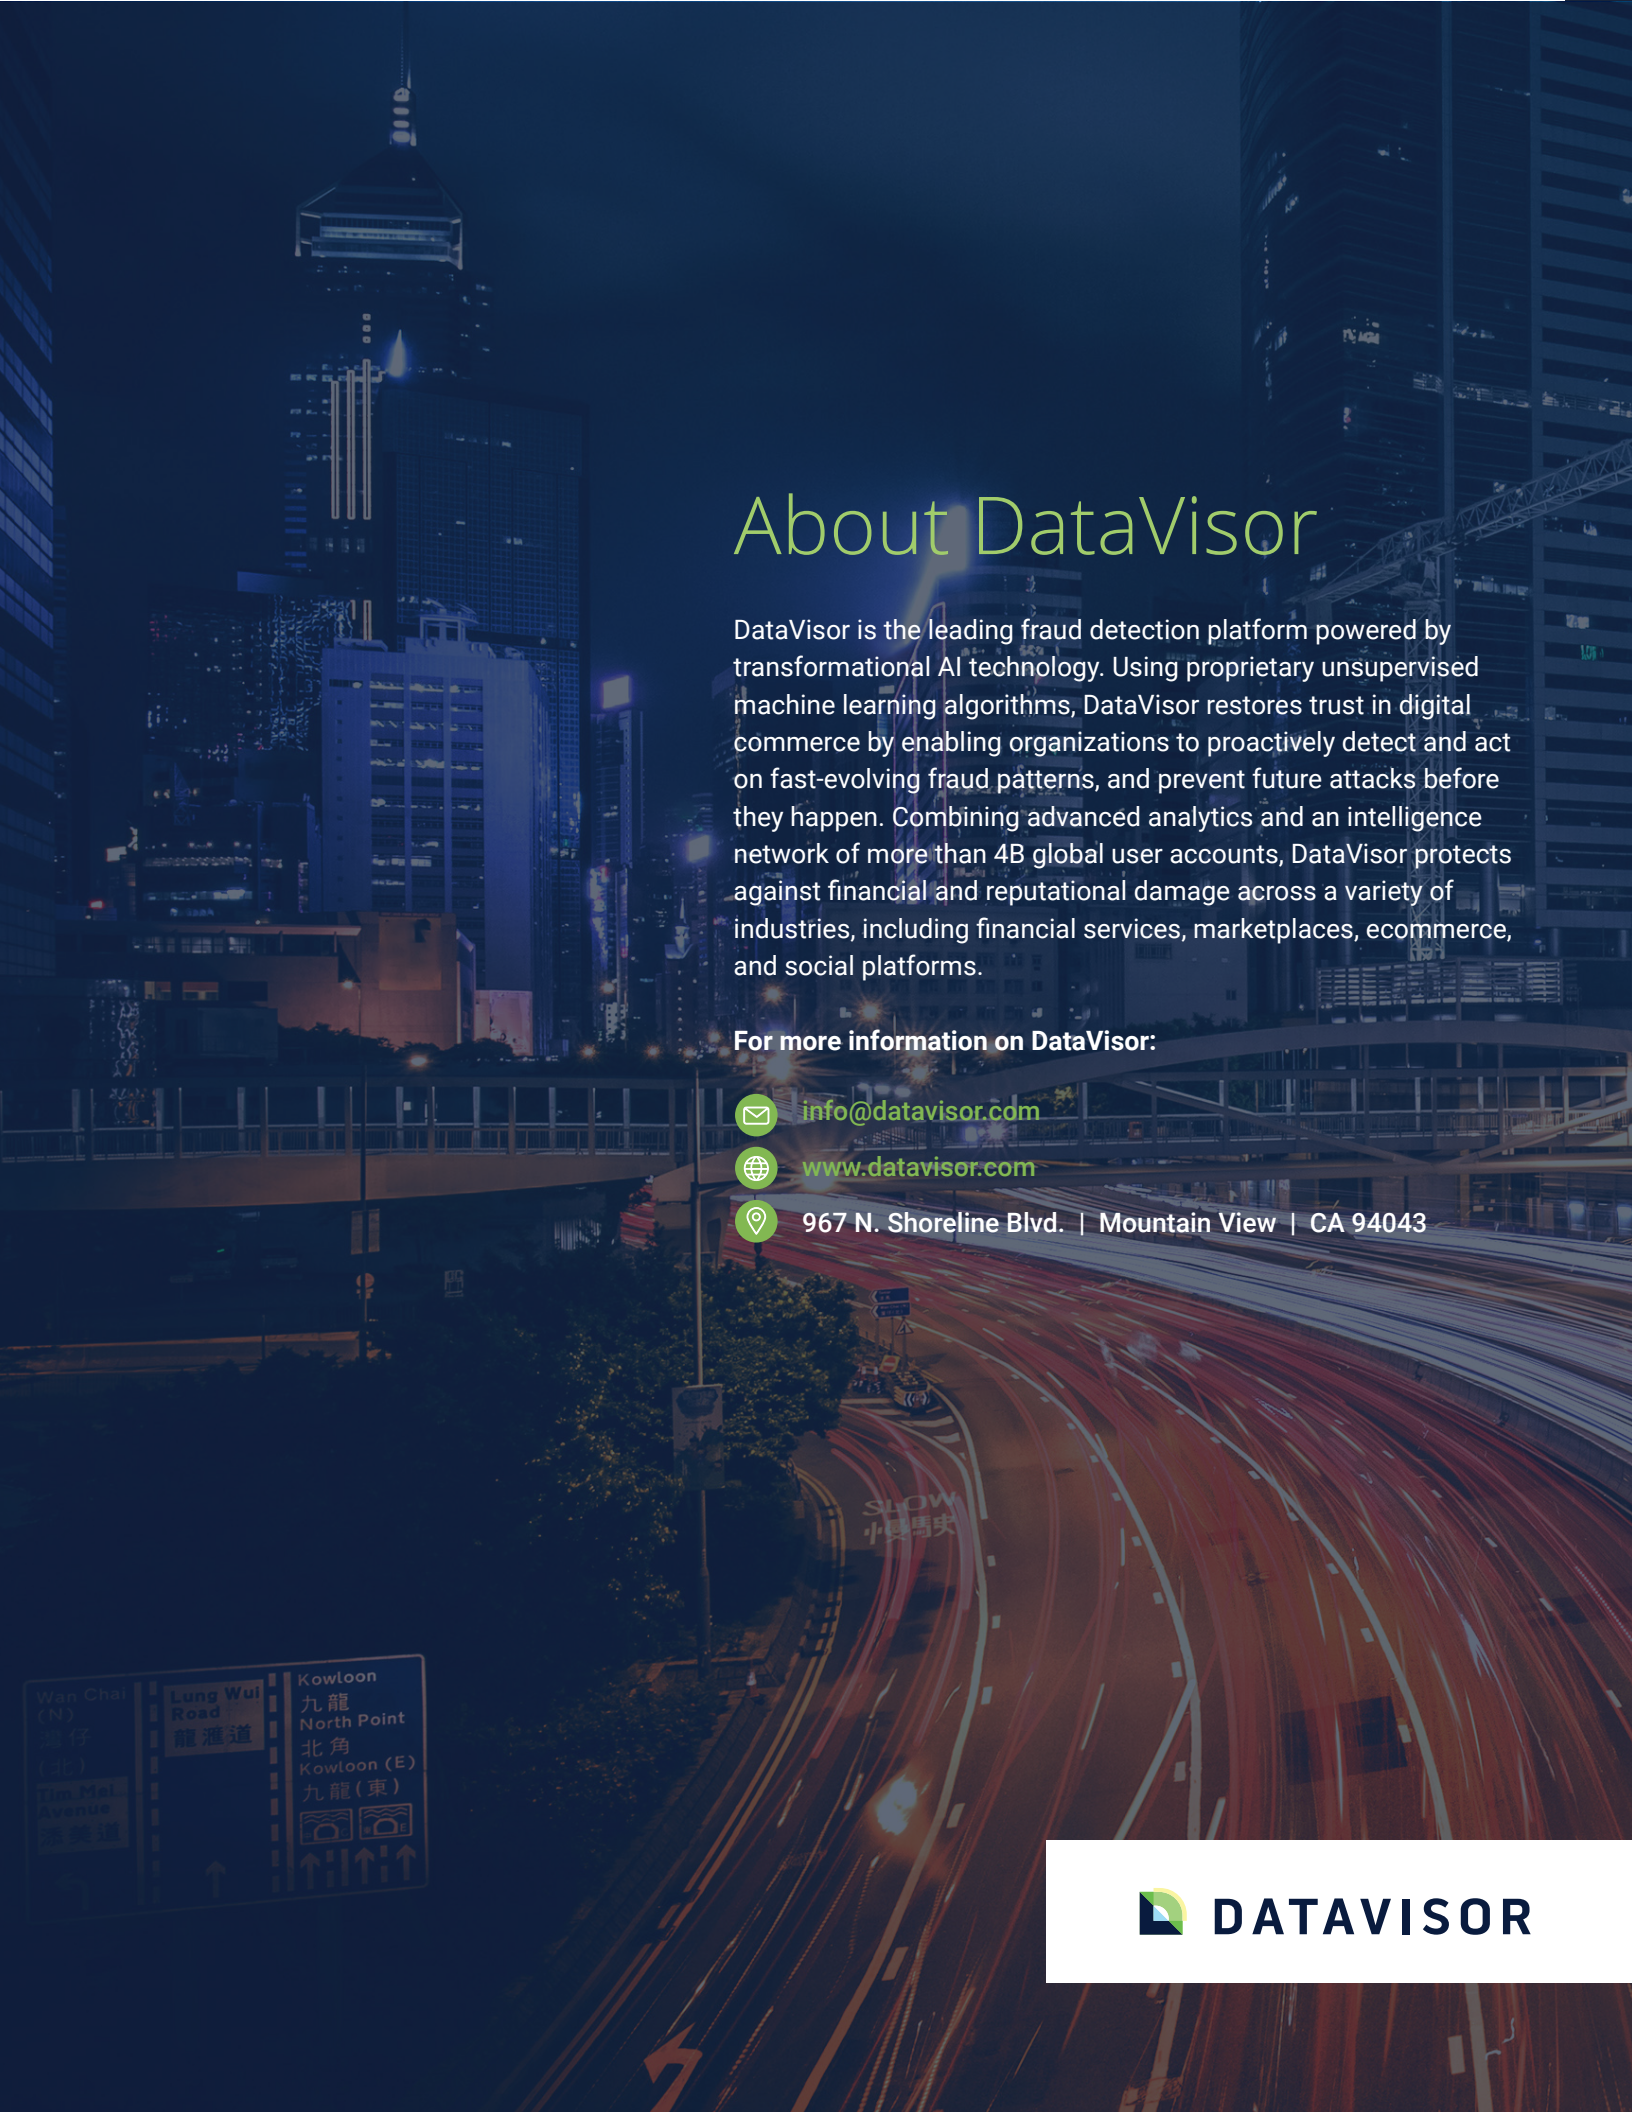

About DataVisor

DataVisor is the leading fraud detection platform powered by transformational AI technology. Using proprietary unsupervised machine learning algorithms, DataVisor restores trust in digital commerce by enabling organizations to proactively detect and act on fast-evolving fraud patterns, and prevent future attacks before they happen. Combining advanced analytics and an intelligence network of more than 4B global user accounts, DataVisor protects against financial and reputational damage across a variety of industries, including financial services, marketplaces, ecommerce, and social platforms.

For more information on DataVisor:

info@datavisor.com

www.datavisor.com

967 N. Shoreline Blvd. | Mountain View | CA 94043

DATAVISOR