



Q1 2020 DIGITAL FRAUD TRENDS

The State of Sophistication in Fraud

A DATAVISOR SPECIAL REPORT

Table of

Contents

EXECUTIVE INSIGHT: THE STATE OF SOPHISTICATION IN FRAUD ATTACKS	3
FOREWORD	5
DEFINING AND UNDERSTANDING ATTACK SOPHISTICATION.....	6
THE DURATION OF MODERN FRAUD ATTACKS.....	8
ATTACK DURATION IN THE CONTEXT OF ATTACK SOPHISTICATION.....	10
ACCOUNT INCUBATION.....	11
ATTACK TECHNIQUES	13
THE INCREASING PREVALENCE OF BOTS.....	14
ATTACK TECHNIQUES BY VERTICAL	15
FRAUD ATTACK GROWTH RATES.....	16
THE LIFETIME OF FRAUD SIGNALS.....	17
FRAUD PREVENTION FOR THE NEW DECADE.....	20
CONCLUSION.....	24

Executive Insight:

The State of Sophistication in Fraud Attacks

We stand today at the beginning of a new decade. Behind us are the lessons of the past, and ahead are the possibilities of the future. We recognize this as a unique opportunity to make meaningful decisions that will shape the future of enterprise growth, customer experience, and risk mitigation for the coming ten years, and beyond. With a vast trove of proprietary data at our disposal, we feel uniquely prepared to offer informed insights as to how organizations can ensure the prosperity of their businesses and the security of their customers.

As we began the work of assembling this annual report, we wanted to derive from our year's worth of data those actionable insights we believed would be most valuable going forward into 2020. We hoped to gather these insights around a central theme, should one prove to emerge. In truth, many themes emerged. However, one stood tall amongst the rest—sophistication.

Digital fraud today is an irrefutably sophisticated enterprise. The range of tools, techniques, and technologies fraudsters have at their disposal, the scale at which bots and automation enable them to operate, and the volumes of stolen and leaked data available to them on any given day, all combine to create a situation in which fraudsters have almost limitless power and possibilities. Add to this all the ongoing new vulnerabilities being exposed through rapid innovation in large sectors such as financial services, healthcare, and insurance, and you have a fraud landscape that is almost indescribably complex.

However, knowledge is power. In the case of modern digital fraud, it is its very sophistication that makes it possible for us to defeat it, because sophistication and complexity by their very nature leave a large footprint—one we can discover, analyze, and act on. No matter how advanced the methods fraudsters may use to obfuscate their efforts, the very act of choreographing thousands of bot-powered fake accounts to launch a coordinated attack means there is a surplus of correlated patterns and cross-account connections that can be exposed with the right solutions in place.

Defeating sophisticated fraud is not easy. The process begins with understanding the concept itself, and this is precisely where our report begins. Over the course of analyzing and digesting this report, you will come to understand what a sophisticated attack looks like, and how sophisticated attacks are built and launched. You will learn the tools and techniques fraudsters use to mount these attacks, and how to identify the signals given off in the process. You will explore what's required to expose and block even the most sophisticated of fraud attacks, and discover how technologies like unsupervised machine learning empower your organization to do so in real time, at scale. Most importantly, you will gain actionable insights into how proactive, comprehensive, AI-powered early detection and prevention strategies can serve to enhance customer experience, and drive business growth. Today, in the new decade, and beyond.

**Yinglian Xie**

Co-Founder and CEO

DataVisor**Fang Yu**

Co-Founder and CTO

DataVisor

Foreword

Sophistication is a signature hallmark of modern fraud, but the meaning and implications of sophistication are not widely understood.

Each day, DataVisor detects more than 30K unique fraud attacks and fraud rings on behalf of our enterprise clients. These attacks differ in myriad ways—by their end goals, the manner in which they're conducted, the tools and techniques used, and more. There are almost infinite differences in the behaviors exhibited by the fraudulent accounts utilized in these attacks, and yet, despite their increasing sophistication, these attacks are both discoverable, and preventable.

Sophistication is a signature hallmark of modern fraud, and yet, while this truism is widely accepted, its meaning and implications are not widely understood.

In this report, we break down exactly what sophistication means when it comes to modern fraud attacks. In doing so, we provide answers to questions such as the following

- ▶ How long do coordinated fraud attacks last?
- ▶ How fast do fraud attacks grow, and how big do they get?
- ▶ What are the differences between attacks with high vs low sophistication?
- ▶ How reliable are fraud signals when it comes to detection and prevention?

Understanding the anatomy and complexity of modern fraud begins by understanding both the commonalities that exist across all attacks, and the singularities that define specific attack types across different platforms and sectors.

Defining and Understanding Attack Sophistication

Low sophistication attacks may seem easier to mount, and accordingly more common, but on any given day, more than half of active attacks can be highly sophisticated.

Implementing the right detection and prevention strategies to meet the challenges of today's fraud attacks requires understanding degrees of attack sophistication. It is important to understand that low sophistication does not equate to easier to block; rather, it simply means different solutions will be required to successfully neutralize the attacks.

The first step in understanding attack sophistication is defining what we mean by the term itself. We can determine the relative sophistication of any given attack based on two primary characteristics:

1. The complexity of the fraudster's attack infrastructure
2. The degree of effort expended to avoid detection

In short, the more complex the infrastructure—and the greater the investment in detection avoidance—the more sophisticated the attack.

► Low Sophistication Attacks

Less sophisticated fraud attacks are more likely to reuse known bad fraud signals, generate noticeably large volumes of malicious activities, and manipulate multiple fraudulent accounts using the same script, such that the accounts have the same profile attributes and behave in the same way.

► **High Sophistication Attacks**

High sophistication fraud attacks are more stealthy, by which we mean they're more deliberately obscured, obfuscated, and disguised, and more subtle in their actions and efforts.

For high sophistication attacks, fraudsters employ multiple tools so that the fraudulent accounts they control can more successfully blend in with other normal users. These fake accounts may have legitimate-looking user profiles with pictures and friends, and they're likely to be older accounts with longer activity histories. They'll often also originate from IP addresses and devices with good reputations.

Sophisticated attacks are generally bigger in scale, and are able to operate under the radar; potentially creating more damage on the online platforms they target.

► **Attack Diversity**

Modern fraudsters are dynamic operators who juggle an array of tools, techniques, and technologies to try and achieve their illicit goals. They often mount simultaneous attacks, and are capable of controlling vast armies of fraudulent and malicious accounts. They attack across several fronts, and combine low and high sophistication attacks depending on their targets and goals. For a business, this represents dizzying complexity, and mandates the deployment of advanced solutions equipped to defend against a full range of attack types.

While low sophistication attacks may seem easier to mount, and while we might conclude they're accordingly more common, the truth, per our research, is that on any given day, more than half of the active attacks we see can be highly sophisticated.




Figure 1: The daily fraction of attacks that are highly sophisticated. Fraudsters are highly dynamic; on some days, more than half of active attacks can be highly sophisticated.

The Duration of Modern Fraud Attacks

*In our research, we observe that attacks on social media sites are the shortest, with a median attack duration of **four days**, while attacks on financial platforms are the longest-running, with a median duration of **31 days**.*

Across all the attacks we've analyzed, the median attack duration is five days. This alone is a telling statistic, as it makes immediately clear the ongoing pressure businesses today face as they attempt to meet scale with scale in the face of increasingly extensive and powerful fraud attacks.

To withstand this kind of pressure requires sophisticated solutions that can act in real time, at big data scale. In addition to speed and scale, these solutions must be adaptive, and able to tailor detection and prevention strategies to specific use cases and platforms.

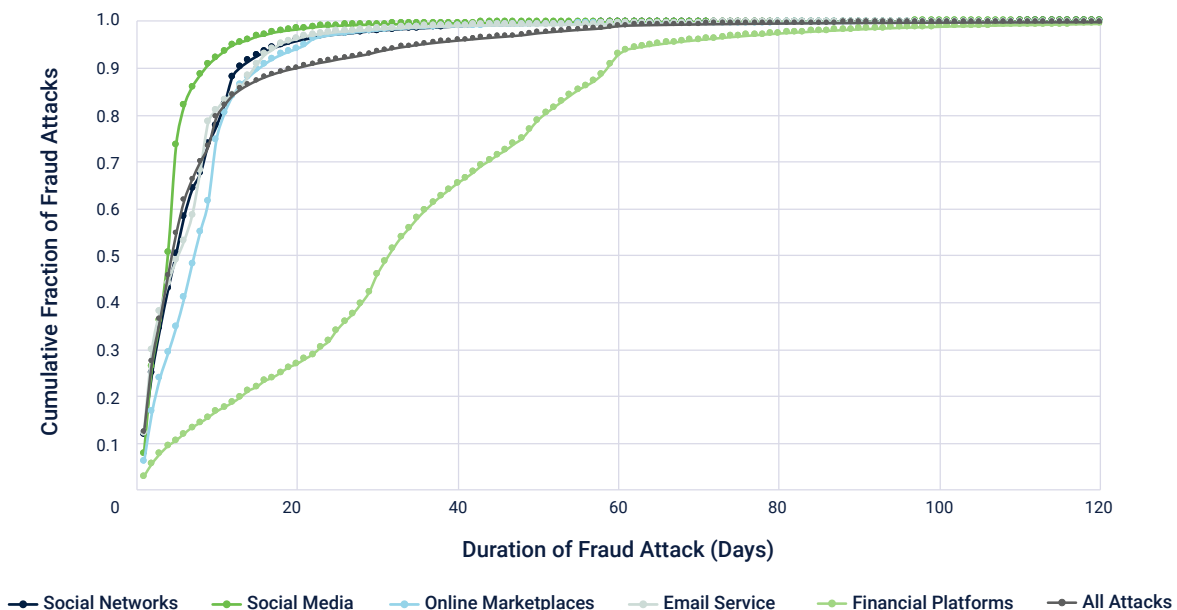


Figure 2: Cumulative distribution of attack duration, for attacks observed on different types of online services.

In our research, we observe that attacks on social media sites are the shortest, with a median attack duration of four days, while attacks on financial platforms are the longest-running, with a median duration of 31 days. The width of this range makes clear that one-size-fits-all solutions can't hope to successfully address the full measure of attack complexity that enterprises today are facing.

Fraudsters on financial platforms maintain long-term operations while cycling through multiple fake accounts or stolen financial information over time. These attacks are typically more stealthy. By contrast, attacks on social networks or social media platforms tend to come in bursts and are comparatively short-lived. The fraudulent accounts used in these attacks are more often than not used for only a single campaign, before being discarded and abandoned soon after.

Attack Duration in the Context of Attack Sophistication

High sophistication attacks can last twice as long as those evidencing low sophistication—15% of low sophistication attacks only last for one day, versus only 8% of high sophistication attacks.

Our research indicates that a majority of fraud attacks last between one to seven days.

No significant difference in attack duration is noticeable for various sophistication levels, though there are 2x as many low sophistication attacks that are short-lived as there are for medium or high sophistication attacks—15% of low sophistication attacks only last for one day, versus 8% of high sophistication attacks.




Figure 3: The distribution of attack duration (i.e., the number of days an attack is active) for varying levels of attack sophistication.

Account Incubation

As many as **87%** of coordinated fraudulent accounts attack within one hour of registration, and up to **92%** do so within **twenty-four hours**.

The question of how quickly fraudulent and malicious accounts get used in an attack is a critical one to answer, as the success of any given fraud strategy will often depend in no small measure on whether solutions in place can see attacks coming in time to block and neutralize them before damage can occur. The challenge of doing so differs depending on whether accounts are used immediately upon creation, or incubated in advance of later deployment.

In our research, we see that a majority of fraudulent and malicious accounts are used in attacks very soon after registration. Between 42%-87% of coordinated fraudulent accounts attack within one hour of registration, and 81%-92% attack within twenty-four hours.




Figure 4: The distribution of incubation time for fraudulent accounts observed on different types of online services.

By comparison, a certain number of accounts will “sleep” for extended periods of time before they’re deployed in a fraud attack. On financial platforms, we see that 11% of fraudulent accounts incubate for more than 30 days before being put to use in an attack, and 5% on social platforms and marketplaces sites.

These are not insignificant numbers, and they serve to highlight the importance of early detection. Account incubation is a two-edged sword. On the one hand, the longer an account stays hidden, the more time a fraudster has to build up a seemingly legitimate digital footprint, and the more legitimate the appearance, the more difficult it is to detect as being fraudulent. On the other hand, wider incubation windows offer businesses more opportunity to identify and block these accounts, provided they’re equipped with advanced early detection capabilities.

Attack Techniques

Fraudsters launch from distributed locations and multiple unique devices to avoid detection, and leverage bot scripts to conduct attacks at massive scale.

There are two imperatives critical to the success of a fraud attack. First, fraudsters need to avoid detection in order to realize their end goal, and second, they need to scale the attack operation to be profitable.

Some of the techniques they use to accomplish these requirements include launching attacks from distributed locations and multiple unique devices (so that each fraudulent account appears like a normal unique user), and leveraging bot scripts to conduct attacks at scale.

Likely due to the prevalence of IP blacklisting and reputation services, launching from distributed IP locations is one of the most common fraud techniques, observed in 46% of all fraud attacks. This is followed by the use of distributed devices; a technique observed in 29% of all fraud attacks. Operating from multiple devices (either via device farms or device emulators) can circumnavigate fingerprinting detection, or otherwise make it difficult to correlate fraudulent activities from the same fraud ring.




Figure 5: The distribution of attack techniques observed in fraudulent attacks.

The Increasing Prevalence of Bots

Location-hopping bots, which originate from IP addresses located in multiple countries or territories, make up 18% of fraud attacks.

Some form of bots are used in virtually all fraud attacks.

► Naive Bots

We observe that “naive” bots, such as those that exhibit strictly periodic activities (e.g., performing an action every five minutes), are relatively rare, likely because they are easily noticed.

► Spiky Bots

By contrast, “spiky” or “bursty” bots that perform many actions within a short time frame before going dormant are much more common—we observe their usage in 22% fraud attacks.

► Location-hopping Bots

Location-hopping bots, which originate from IP addresses located in multiple countries or territories, make up around 18% of fraud attacks. These multinational attacks either consists of multiple groups of fraudsters, or have established network infrastructure (private proxies or VPNs) in multiple locations.

Attack Techniques by Vertical

Data center or proxies are more often used in attacks on social platforms, compared to financial services, likely due to the latter deploying more strict controls related to IP reputation.

While launching attacks from distributed IP addresses appears to be popular among fraudsters regardless of the type of online platform, other attack techniques have more variability.

Social platforms have the highest fraction of spiky, bursty bot attacks, which are commonly used to perform large-scale content abuse.

Datacenter or proxies are more often used in attacks on social platforms, compared to financial services, likely due to the latter deploying more strict controls related to IP reputation.




Figure 6: The distribution of attack techniques observed in fraudulent attacks on different types of online services.

Fraud Attack Growth Rates

One-third of attacks grow at a slow pace—under 50% of their original size—but 15% of attacks are fast-growing; doubling their size or more overnight.

Fraud attacks grow over time by gaining new fraudulent accounts, either from registering fake accounts, or compromising existing accounts.

Approximately 52% of fraud attacks are “static”; meaning, they do not gain more fraudulent users over time, but acquire them in one shot.

To measure the rate at which fraud attacks grow, we calculate the percentage increase of the fraud ring size on each day compared to the previous day, and take the maximum value observed during an attack’s lifetime.

Roughly one-third of attacks grow at a relatively slow pace—under 50% of their original size—but 15% of attacks are fast-growing; doubling their size or more overnight.

The figure shows the distribution of the maximum daily percentage increase in fraud ring size.



Figure 7: The distribution of growth rate in observed fraud attacks, calculated as the maximum percentage increase of the fraud ring size on each day compared to the previous day.

The Lifetime of Fraud Signals

In 2019, we observed 104K unique IP fraud signals across 1.57 million fraud attacks. The median lifetime of IP fraud signals across all attacks is 4 days.

To understand the importance of fraud signals for fraud prevention, it is first necessary to be clear what we are referring to. We refer to “fraud signals” as those characteristics that are shared by the majority of fraudulent accounts in an attack. Fraud signals may be IP addresses, device types, user-agent strings, nicknames, or common pieces of content.

A fraudulent account may originate from many different IP addresses or switch between different endpoints or devices over the course of the attack. However, when they are controlled by the same fraud ring, there are always giveaways that tie those accounts together to the same operation. By spotting and identifying these signals, we are able to close in on a specific fraud ring, which is how we ultimately neutralize attacks.

► IP Fraud Signals

IP fraud signals play an important role in our approach to fraud prevention, as all attacks—regardless of the online platform—need to originate from a network location. We can measure the use of these fraud signals across time, across fraud attacks, and across online platforms. Over the course of 2019, we observed 104K unique IP fraud signals across 1.57 million fraud attacks.

The median lifetime of IP fraud signals across all attacks is 4 days. This means that a fraud attack will only utilize the same IP address for 4 days, after which they move on to a new address.




Figure 8: The median lifetime of IP fraud signals on different types of online services.

IP fraud signals used in attacks with higher sophistication have a slightly longer lifetime. This correlates with our earlier observation that high-sophistication attacks have longer durations.

By operating longer and stealthier, these attacks can potentially cause much more damage on the platform.

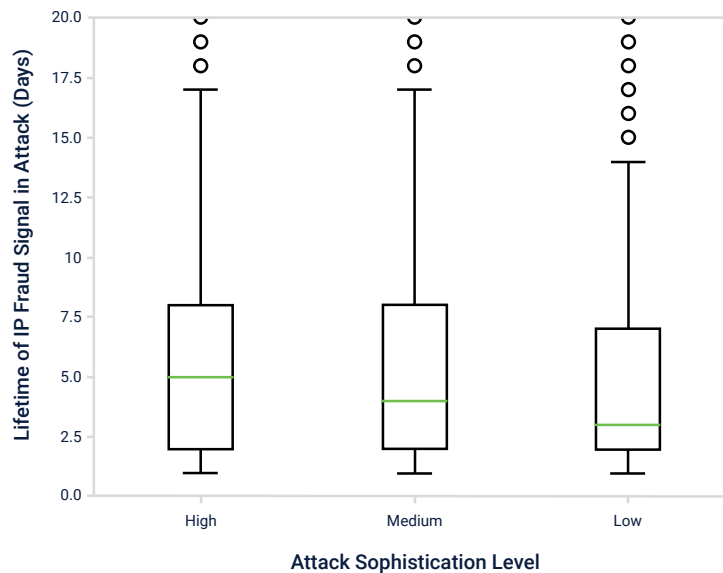


Figure 9: The distribution of the lifetime of IP fraud signals observed in attacks of varying levels of sophistication.

We also observe a high “reuse” rate of IP fraud signals, especially among attacks targeting social platforms or online marketplaces. On average, an IP fraud signal is used in 24 attacks on social networks, and 23 on marketplaces—around 6x the rate of those found on financial platforms. Financial attackers are more sophisticated, and take measures to avoid reusing signals with poor reputation.

Among IP signals, 26% are used in fraud attacks across multiple online services. This shows that fraud rings are likely to operate across online services, and also serves to highlight the prevalence of infrastructure reuse, which is often due to fraudsters purchasing the same hosting services offered in the fraud-as-a-service underground economy.




Figure 10: The number of attacks associated with each IP fraud signal, for attacks observed on different types of online services. This shows the frequency at which IP fraud signals are reused by fraudsters.

► **Email Fraud Signals**

Email-based fraud signals are also valuable for proactive fraud prevention. In a previous DataVisor research report, we observed, in a single quarter, more than 490K IP addresses, 17K user-agent strings, and 1.5K email domains that could be confirmed as fraud signals. Out of these, 36% were only active for less than one day, and 64% were only active for less than one week. IP addresses were the most volatile.

A closer look at the email domain fraud signals showed that around one-fifth were registered within the last two years. Registering private domains allows fraudsters to create email accounts en masse, bypassing phone verification, CAPTCHA, and other authentication methods often required for public email services.

Fraud Prevention for the New Decade

Across today's fraud landscape, we can see a vast array of sophisticated attack types that run the gamut from application and transaction fraud to money laundering and identity theft. Yet while there is indeed great attack diversity, one phenomenon is a near-constant across so many different instances of modern digital fraud—data breaches.

Every sophisticated fraud attack has a kind of timeline—from the first moment the attack is conceptualized, to the penultimate point when it is either monetized or neutralized—and in countless use cases, we find data breaches playing an essential role. The personal information that gets exposed in a data breach all too often becomes the raw material that makes future attacks possible. Synthetic and fake identities are built using information bought and sold on the dark web, and those identities are, in turn, used to mass register accounts that subsequently become soldiers in a fraudster's army.

Money laundering—and its associated “mule” activity—is a massive global concern. According to the [United Nations Office on Drugs and Crime](#), “the estimated amount of money laundered globally in one year is 2 - 5% of global GDP, or \$800 billion - \$2 trillion in current US dollars.” And, per a recent [LexisNexis report](#), “annual anti-money laundering (AML) compliance costs for the United States and Canadian financial institutions totaled \$31.5 billion in our recent study.”

Money laundering can be a highly sophisticated process, and data breaches play a central role. In an exclusive conversation with DataVisor for this report, Julie Conroy (Research Director, the Aite Group), offered a look ahead into the new decade. While she sees ongoing challenges with data breaches, she also finds cause for hope:

“ We definitely won't see breach rates go down. In fact, it will likely be quite the contrary based on the *recent Verizon Data Breach report, which shows that full PCI compliance dropped from 52% in the 2018 report to 36.7% worldwide for the 2019 report*. The organized crime rings behind this fraud are well funded and nimble, whereas their targets have to build business cases and wait in IT queues to deploy compensating controls.

Amid the doom and gloom, however, one hopeful sign that I'm seeing out there is the growing appetite among banks to take action on the mule networks that enable so much financial fraud and money laundering. Traditionally it's been difficult to get funding to address muling, since the mule activity in and of itself does not typically cause a direct loss to the FIs where the accounts are housed (unless it's on-us activity.) Many FIs have long recognized, however, that mule activity is a key enabler of financial fraud, and many large FIs are now taking action, either individually or collaboratively. If we can put a dent in the mule networks, then we will certainly make it more difficult for organized crime to monetize the results of all the breaches and remove the funds from the system. ”

—JULIE CONROY, RESEARCH DIRECTOR WITH THE AITE GROUP

Richard Cooney (Director of Fraud Strategy, Axxess Financial), also spoke with DataVisor for this report. His comments make clear the role of stolen data in the modern digital fraud ecosystem, and serve to shine a light on the global nature of the challenges ahead:

“**W**hen it comes to fraudsters, more is always better. Low prices for credit card and identity data continue to make it highly profitable for fraudsters to target the business sector using velocity attacks. Hence, business losses remain high. This is mainly due to the sheer volume of data available for sale on the dark web.

The international law enforcement community lacks the resources to exert real pressure on those who buy and sell products on the dark web. We see the occasional arrest sensationalized by the press, but these arrests do little to impact activity on the dark web. There are too many countries that continue to provide a safe haven for cybercriminals.

I believe market forces will continue to drive growth in hacking activity. The power to stop this growth and the related crimes remains with businesses themselves. Businesses must continue to modernize their systems, and in doing so, build in the appropriate data safeguards and system controls to identify and deter hackers. Businesses susceptible to credit card and identity fraud must continue to innovate. They must invest in new tools, build layered controls, close gaps, and develop new strategies to make the cost to defeat their controls too high for the fraudster.”

—RICHARD COONEY, DIRECTOR OF FRAUD STRATEGY AXCESS FINANCIAL

Given the sophistication and complexity of global digital fraud, it's no surprise that Richard advocates modernization and innovation when it comes to fraud prevention in the coming decade. Steven D'Alfonso (Research Director, Worldwide Compliance, Fraud and Risk, IDC Insights)—who also spoke with DataVisor for this report—echoes Richard's sentiments about the critical importance of transformational technologies such as artificial intelligence when it comes to fighting sophisticated modern fraud—particularly in the finance sector:

“ **A**dvanced fraud analytics using AI technologies are proving to be effective in identifying fraud and preventing losses. As the industry increasingly moves into real-time payments, it is necessary for banks to upgrade fraud management tools to identify real-time fraud attacks and detect changes in fraudsters' methods of operation. Financial institutions that are slow to implement newer fraud detection technologies will become bigger targets as fraud activity is displaced from institutions that have implemented them. ”

—STEVEN D'ALFONSO, RESEARCH DIRECTOR, WORLDWIDE COMPLIANCE,
FRAUD AND RISK, IDC INSIGHTS

As we look ahead into 2020, it is clear that challenges loom. What is also clear, however, is that we have the means to meet and overcome these challenges, using the power of advanced, AI-powered fraud solutions.

Conclusion

Over the course of this report, we have established several key points about the character and make-up of modern fraud attacks. First and foremost, we have shown that fraud attacks today are capable of evidencing extremely high sophistication. We have defined sophistication, and made clear that these attacks are potentially massive in scale, that they can last for extended periods of time, and that they are increasingly difficult to detect due to advanced obfuscation techniques and strategies.

We have also illustrated techniques that fraudsters use to grow, extend, and obscure their attacks, and we have discussed at length the signals given off by fraudsters that enable our solutions to identify and neutralize their attacks.


The most important takeaway from this report should be the clear need for proactive detection strategies that can analyze huge volumes of data in real time, without the need for labels, so as to isolate those signals that indicate coordinated and connected activity. Without these capabilities, it is virtually impossible to spot high-sophistication fraud attacks in time to prevent extensive damage from occurring.




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**