CEQUENCE® SECURITY

API SECURITY THREAT REPORT

Bots and Automated Attacks **Explode**

Contents

INTRO

2021 Attack Observations

Looking back on 2021, we found that attackers and legitimate businesses have similar characteristics in that both groups are opportunistic, making the best out of what they have. Businesses were forced to move online to survive due to the pandemic. Attackers followed, taking advantage of the new, increased volume of targets, most of which were application programming interfaces (APIs).

Traditionally, bots have targeted web applications as that is where they could find the most success in gaining access to accounts or creating attacks that focused on achieving their end goal, be it account credentials, or buying up highly prized inventory.

During the last two years, organizations have moved away from a web-based application infrastructure to one that is based on APIs. Moving to APIs allows organizations to address several different development goals: support for mobile, the move to the cloud, faster release cycles and greater ecosystem connectivity.

The move to APIs also enables a shift in attack targets from web channels with headless browsers to APIs with nothing more than command line tools that make perfect computer-to-computer requests at computer speed. The numbers prove that both developers and attackers have made the shift — of the 21.1 billion transactions analyzed by Cequence Security in the last half of 2021, 14 billion (70%) APIs.

Looking beyond the sheer volume, attackers were very active in their retooling efforts when presented with roadblocks or restrictions on what they were trying to achieve. The time it takes to retool is often observed as a way to gauge how much attention is being paid to the attack as well as the sophistication level of the attacker. When we see rapid retooling to avoid roadblocks , we know the attacker is watching and is capable of making changes.

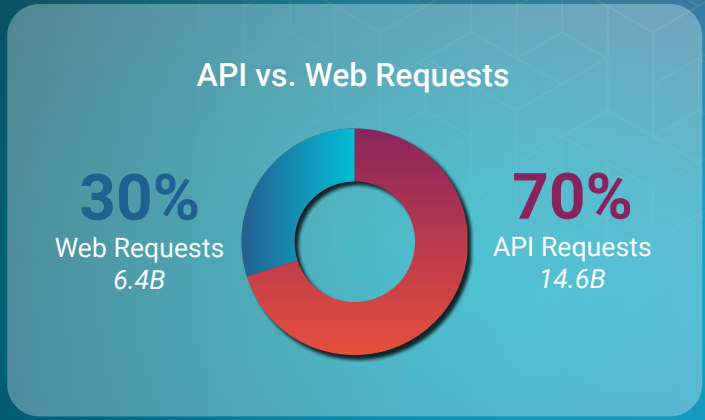

Image 1: *June to December 2021application request analysis.*

Looking back at 2021, the CQ Prime Threat Research team discovered that attackers prioritized three key techniques for wreaking havoc in industries from retail to banking: fraud, shopping bots and account takeovers.

ATTACK TREND ONE

Fraud Comes in Many Forms

Fraud of some type is the most common end-goal of the automated attacks prevented by the Cequence platform. The three examples discussed below include retail gift card fraud on a shadow API, manual fill loan application fraud and payment fraud. All of these examples required a high level of attacker understanding and sophistication to achieve success.

GIFT CARD ATTACKS

LOAN FRAUD

PAYMENT FRAUD

High-Volume Gift Card Fraud Attacks

2021 saw a big uptick in gift card fraud across all of our retail customers. We uncovered a scheme that had gone undetected by the incumbent solution and had resulted in roughly $200K in losses over time. Another set of attacks that stood out was a coordinated series of high-volume attacks in late July that saw retail customers get hit with a 28X increase in malicious requests, averaging 700K per day. The unique characteristic was that the attackers were using low-cost, no-log commercial VPN offerings as their attack infrastructure, allowing them to mask their identity and location. These VPN connectivity solutions were mapped back to several vendors including IPVanish[1], Highwinds Network Group, StackPath Data Center, Netprotect, Reliable Hosting, Inc. and Overplay.com. Attacks were primarily account takeovers with the end goal of gift card and loyalty points fraud against retail customers.

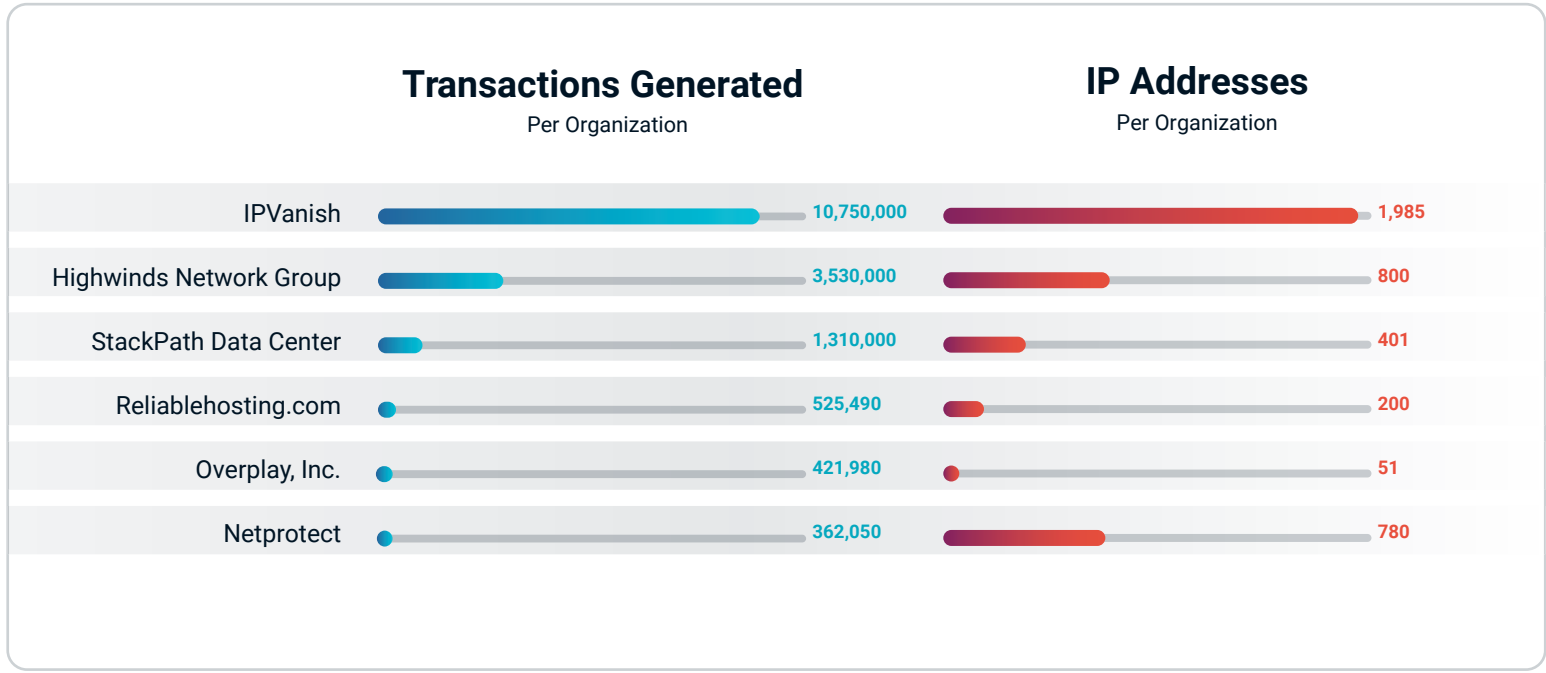

Image 2: *Transaction and IP address distribution across different organizations.*

High-Volume Gift Card Fraud Attacks

The legitimate use of VPN services makes retailers' "block or not?" security decision difficult. Prior to the spike observed on July 21, 2021, traffic from these organizations were legitimate users, or limited to periodic, relatively targeted (and mitigated) attack bursts. As shown in image 3, on July 21, threat actors drove a 57X traffic spike to 1.44M daily transactions reaching a high of 1.87M per day on August 4.

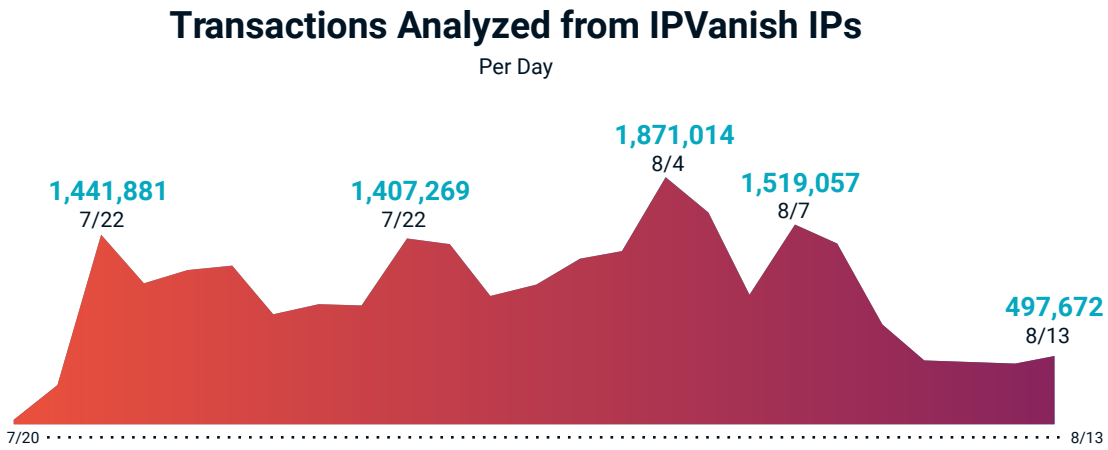

Image 3: *Daily transactions from IPVanish and related organizations.*

Traffic flowing through these organizations' IP addresses targeted as many as 20 different customers across multiple industries. However, more than 99% were targeted retailers with popular gift card and loyalty programs. Continuing the trend observed over the past year, threat actors continue to target APIs as opposed to web apps, with 90% of all attack traffic funneled through APIs.

Industry	Transactions Analyzed	% of Total
Retail / E-Tailer / E-Commerce	16,098,621	99.6%
Romance / Relationship	35,425	0.2%
Financial Services	18,410	0.1%
Digital Media	3,689	0.0%

High-Volume Gift Card Fraud Attacks

Gift Card Attack Characteristics Observed

While gift card attacks were seen in multiple industries, the focus was on retailers specifically to execute ATO and gift card fraud as described below.

- **Scrape and sell gift card abuse:** The attackers are using a bot to commit account takeover, log into the compromised accounts and request gift card balance information from the profile APIs. This information would then be resold later on the market for future abuse.

- **Steal and purchase gift card abuse:** The attackers used a bot to commit account takeover, using the compromised account to purchase goods, paying for them with the value extracted from the gift card. In some cases, this appeared as attackers using loyalty points to then purchase more gift cards, which were then redeemed for goods.

- **Shopping bot gift card abuse:** Attackers used a bot to execute a shopping journey faster than a competing human, and then during the checkout phase, attempted to enumerate possible gift card numbers against a separate API, applying any valid numbers to the checkout.

Evasive Techniques Employed

- **Masking themselves as a Googlebot:** In some cases, the threat actors tried to mask themselves as a Googlebot by modifying the user agent string to Googlebot/2.1 (http://www.google.com/bot.html). Security analysts should review Googlebot Documentation and observe the characteristics of legitimate Googlebot traffic. For instance, is it coming from Google Cloud, Google Partners, Google inc., etc. by DNS ISP or Organization name records? Analysts should also ensure that Googlebot user agent strings are not whitelisted by default.

- **Using nurtured or fake accounts:** In other cases, attackers used a set of fake accounts that had been harvested over time. As these accounts were often created by the attackers themselves, the patterns don't look like account takeovers, as the attackers always log in successfully. These nurtured accounts can later be sold, or used as mule accounts to execute any purchases/transactions on the site.

Low and Slow Attacks Lead to Loan Fraud

Automation is often equated with high-volume or widely distributed attacks. In this example of loan fraud, the automation was mostly human, and was intentionally execute in a low and slow manner – possibly to mask the attacker actions. To uncover the malicious activity, data samples to find patterns in the application requests were collected over several days to accurately identify where the attacks were taking place. The CQ Prime Threat Research Team has labeled this attack as being very sophisticated, involving more than one individual and was likely manual, given its low and slow nature. The perpetrators consistently adjusted several signals during this attack, requiring diligence to defend against it.

During the early stages of the attack, the CQ Prime Threat Research Team identified roughly 3,000 email addresses distributed across as many as 10 different IP addresses per loan application. Moving IP addresses once or twice in a session is not uncommon, but 10 times is a pattern indicative of suspicious behavior, confirmed by additional analysis.

The attackers were trying to minimize the work on their end by leveraging a standard Gmail account feature – the ability to use subaccounts. Rather than create many individual email accounts (each of which require the attackers to login to check status), they created subaccounts using dots [.] in the primary email account. (Example: Primary email: abc123@gmail[.]com, subaccount email abc.123@gmail[.]com and a.bc123@gmail[.]com.) This gave the perpetrators the ability to more easily submit multiple loan applications to increase chance of success (the law of numbers), but only one is used to consolidate status checking and communication.

The discovery of this pattern combined with IP address data allowed the Cequence team to create policies and block the activity. Yet another tactic used by the attackers was large-scale cookie manipulation, likely executed by a large team of humans running limited automation to generate this (low) volume of traffic. Discovery of this technique in action uncovered roughly 45,000 fraudulent loan applications, which were canceled immediately.

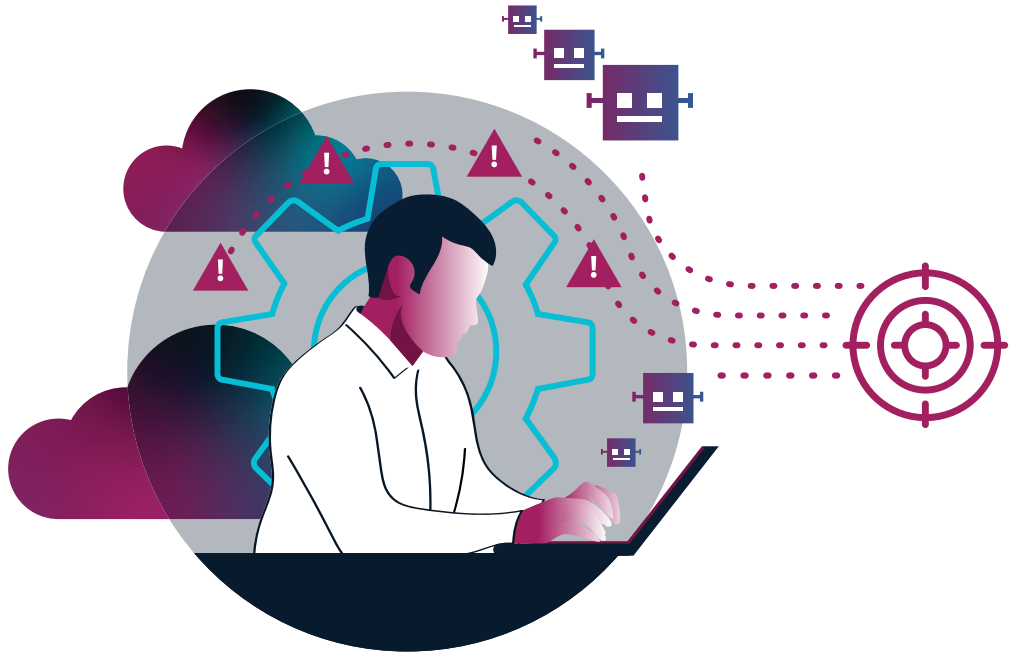

(Too) Consistent Timing Uncovers Sophisticated Payment Fraud

Payment fraud takes many forms, but this 2021 case was narrowly targeted and highly automated, going undetected by the previous bot prevention solution for a year or more. The attackers targeted an API, making regularly spaced payment authorization calls from more than 20,000 phone numbers, all emanating from three zip codes and representing 5.1% of the company's overall payment traffic.

Unlike some of the other attacks where the sheer volume of traffic was a giveaway, this attack was uncovered by correlating several different patterns. The CQ Prime Threat Research team dug into the origin point of the traffic (same three zip codes), when the requests were made (same timeframe each month), the phone numbers used and the timing of the billing request (too perfectly timed). The result of the analysis uncovered payment fraud that was quickly shut down via a blocking policy.

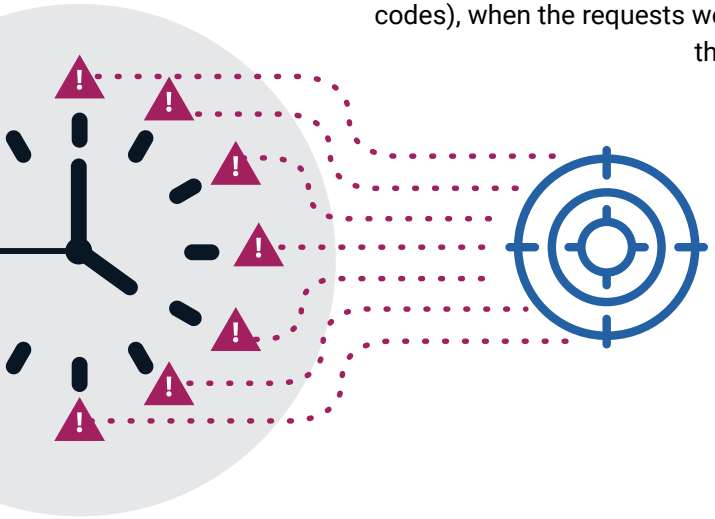

ATTACK TREND TWO

Shopping Bots Get More Sophisticated

Shopping bots are often viewed as the most technically sophisticated bots, and are driven primarily by the revenue opportunity in secondary markets for high demand sneakers, game consoles, luxury items, and graphics cards. As an example, the secondary sneaker market is forecast by Cowen Financial to be a $30B market by 2030, making it a worthwhile investment to try and improve the botting technology. Accordingly, Cequence saw significant investment in the form of increased commercialization that effectively made bots available to the masses.

ENTER
BOTS-AS-A-SERVICE

SHOPPING BOTS
& THE BOTTOM LINE

Enter Bots-as-a-Service

At one time, a bot manager would need to find the tools, scripts, credentials and infrastructure to assemble an attack. Now, bots-as-a-service allow you to buy, rent, and subscribe to a bot and then use it to acquire the (high demand) item of your choosing. The bot commercialization trend is having the largest impact on hype sales where limited-quantity, high-demand items are launched and the retailer is slammed with legitimate users and automated-bot-generated requests. One of our retail customers performs 3-hour hype sales on a regular basis.

At any one given time, daily traffic volume for this customer roughly 1M transactions per hour or 3M over the typical product launch. Bots hit this customer so hard that every launch they execute sees traffic spikes ranging from 12X to 43X above normal with up to 86% of the transactions being malicious.

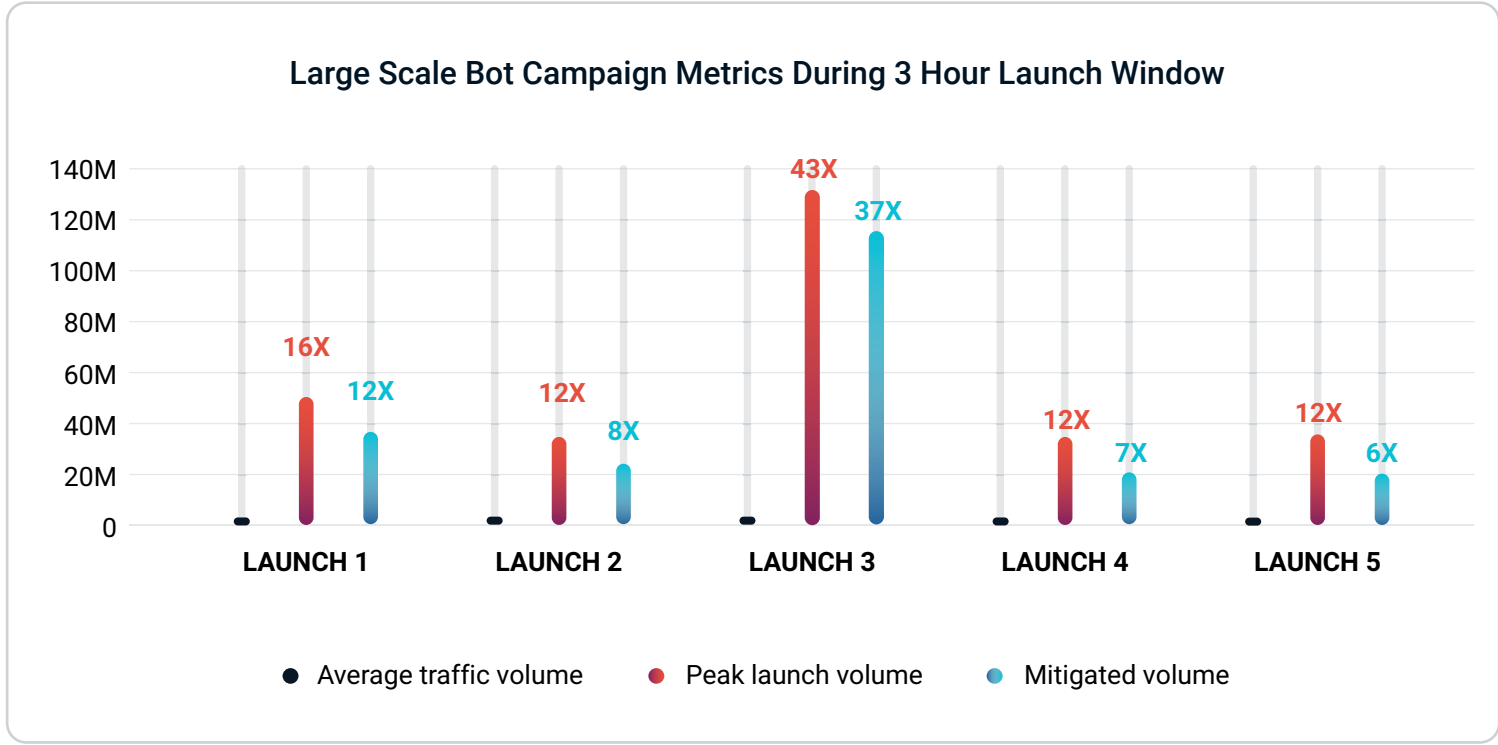

Image 4: *Malicious traffic spikes observed during hype sales events.*

Enter Bots-as-a-Service

Shopping bots are not only more readily available, they are also highly sophisticated – when there is money (recall the $30B projected market) to be made, attackers will spend as long as it takes to maximize their efforts. While retailers are preparing to execute a successful launch, bot managers are doing the same thing – testing their tools, building out their infrastructure, probing their targets for new avenues of access. A few of the preparation steps they take are outlined below.

- **Inventory tracking:** Oftentimes, inventory is loaded prior to the launch and bot managers will use the check inventory API to find out when the item is available. Once they have the information in hand, they will use automation to log in to their fake accounts and pre-load their shopping cart, using techniques discussed below.

- **Account farming:** Bot managers need many fake accounts to help ensure their success. Two newer techniques observed in 2021 were Gmail farming and domain parking. Gmail farming is where the bot manager will use manual and automated efforts to painstakingly validate a Gmail account in preparation for use. Once validated, the one Gmail account can be used to create several subaccounts, as indicated above. Domain parking is a technique where the bot manager registers a fake domain, complete with unlimited email accounts and mail forwarding.

- **Cart build-up, cart farming and product switching:** Bot managers have figured out that many retail shopping carts do not expire. To take advantage of this, bot managers will load many shopping carts associated with their fake accounts with unwanted or very low-cost items. As the drop occurs, they will add the desired product to the cart for more rapid execution.

- **Third-party one-click purchase APIs:** Another technique attackers use to ensure success is finding one-click pay APIs commonly used by ApplePay, Google Pay and PayPal Express, which allow them to execute their purchase more quickly.

Shopping Bots and the Bottom Line

If a shopping bot results in a sale, should the retailer even care? The answer is yes. Shopping bots may generate revenue as the sale intended, but the impact of the attack(s) is far-reaching, touching nearly every department. Fraud teams get overloaded with account validation/reset efforts. IT teams become burdened with traffic spikes beyond their budgets, down websites and mobile apps. PR teams deal with frustrated buyers' social media outcry and marketing departments are faced with faulty KPI statistics. As a proof point, the November 2020, State of Online Fraud and Bot Management Report by Forrester research, showed that 63% of the respondents felt that shopping bots (called inventory hoarding) had a negative impact on their bottom line that ranged from 3%-10%.

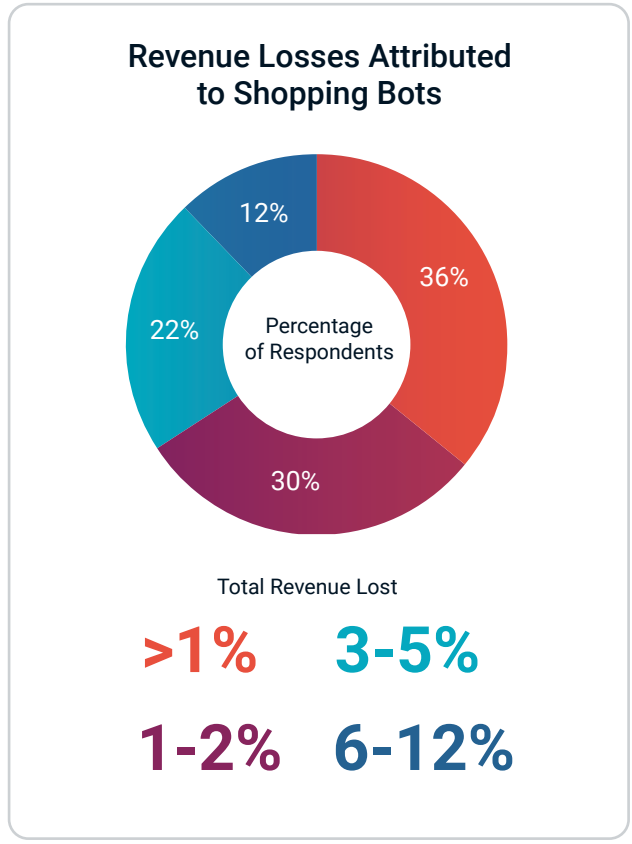

Image 5: *Revenue losses attributed to shopping bots.*

ATTACK TREND THREE

The Account Takeover Cat-and-Mouse Game

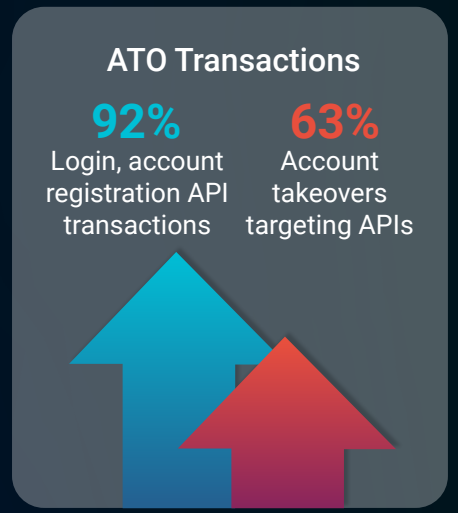

Image 6: *June to December 2021 ATO transaction analysis*

The analysis of more than 21 billion transactions from June to December 2021 showed account login and registration API transactions increased by 92% to more than 850 million. Highlighting the fact that attackers love APIs just as much as developers, that same dataset showed account takeover (ATO) attacks on login APIs increased by 62%. The initial impact of an ATO on an end-user is to panic. You have received a password reset notification from your favorite retailer/social media/ financial institution because your account has been compromised. Being a victim of Account Takeover isn't very fun and causes one to want to stop doing business with the organization the account is for.

ATTACK TREND THREE: THE ACCOUNT TAKEOVER CAT-AND-MOUSE GAME

ATO techniques have expanded beyond the commonly used hot and heavy, high-volume credential stuffing to include methodical low and slow attacks using specific usernames and passwords. This pattern is best exemplified in a customer that has a significant social aspect (i.e., reviews, recommendations, etc.) to their platform.

For them, account takeovers have become a persistent problem where the goal is to not necessarily steal from the compromised account, but to use it to amplify positive or negative information. Between mid-July and late September 2021, this customer experienced four distinct attack-defend cycles outlined below.

- **Phase 1**: ATO mitigation was enabled during the first week of July resulting in an immediate impact. Bots disappeared – but only for about a week.

- **Phase 2:** In this phase, bots returned in full force beginning in late July, continuing for nearly 2 months with high volume attacks consuming up to 80% of all login traffic. The bots use thousands of clean, residential proxies to distribute and anonymize the attack, striving for a 1:1 ratio of IP to transaction. The attackers also continuously retooled their efforts every-other-day using techniques such as randomized user agent string rotation, valid header and body value replay and farming a broad array of Google Analytics cookies.

- **Phase 3:** In yet another attempt to hide in plain sight, during this phase, attackers shifted from high-volume to low and slow, mixing their evasive techniques. Attackers were seen reverse-engineering good browser fingerprints and farming legitimate cookie profiles (both site specific and common web tools like google analytics). For a period of more than 3 weeks, bot activity was low and slow, never exceeding 20% of overall traffic.

- **Phase 4:** Bot activity has dropped to nearly an all-time low, for now.

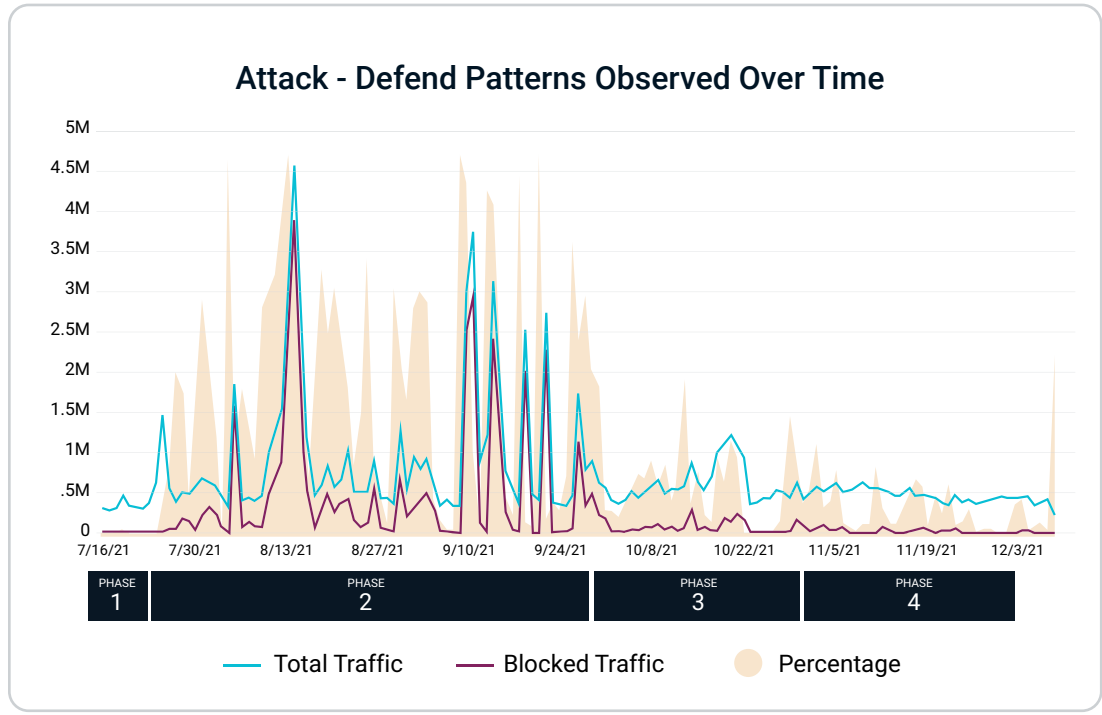

Image 7: *ATO attack - defend cycles observed between July and December 2021.*

The patterns observed here have been seen previously in one form or another in other customer environments. Bots go quiet for a time period and they return with a vengeance. Monitoring bot forums confirms that bot managers often collaborate, sharing ideas, probing for unprotected vectors, like a deprecated API, all in preparation for the next attack. A successful defense requires continued vigilance, monitoring all types of endpoints – web, API and mobile and collaboration between your peers and with your protection provider.

Summary

Looking back at some of the more interesting bot attacks in the past year clearly demonstrates that bots have shifted their focus to APIs and have simultaneously increased their levels of sophistication. Moving into 2022 will bring increased automation against exposed APIs and the use of AI/ML by attackers to successfully bypass commercial defenses.

1 The IP address data collected during these attacks was shared with J2 Global, the IPVanish parent company who took action against the malicious users.

Protect Your APIs While Empowering Your Developers with Cequence Security

Schedule Your Cequence API Security Platform Demo at cequence.ai/demo

CEQUENCE® SECURITY

100 S. Murphy Avenue, Suite 300, Sunnyvale, CA 94086 | 1-650-437-6338 | info@cequence.ai | www.cequence.ai