Fighting ATO Attacks: Protect Your Brand & Customers

A guide for eCommerce merchants

April 2018

Introduction

Online fraud is becoming more and more sophisticated, as cybercriminals try and keep a step ahead of fraud solutions and tools. One of the results of this arms race is the recent surge in ATO – account takeover – attacks, a form of fraud which is particularly difficult to detect. In 2017 ATO led to $5.1 billion in losses, a staggering 122% increase over the $2.3 billion lost in 2016.

Not only are ATO attacks tough to spot, they can also cause harm that goes beyond just stolen goods and chargebacks. Customers often leave credit card details saved in their store accounts, trusting merchants to guard them. In the event of a data breach, customers are left to deal with the fallout of having their personally identifiable information–PII– stolen. This could entail cancelling credit cards or dealing with identity theft. It can even be a struggle simply to regain control of their account: once fraudsters have accessed an account, they can lock the owner out by changing the security questions and passwords. In 2017, ATO victims spent an average of 15 hours resolving the fraud.

The damage for merchants is at least as grave. ATOs reflect very poorly on their brand, create a breach of trust with loyal customers, and can potentially lead to the loss of the entire lifetime value of affected customers.

In this guide, we'll explain ATO attacks: how fraudsters get the credentials they need to access accounts, and their modus operandi once they're in. We'll also explain how to protect customer data, and your own products, from sophisticated ATO fraud. Finally, we'll provide tips on creating a verification process that keeps bad actors out, without causing unnecessary friction to your good customers.

How do fraudsters get credentials to carry out ATO attacks?

The story of an ATO starts with a data breach. There are no shortage of examples of huge hacks in the news; Equifax, Yahoo, Target, and Adult Friend Finder have all been victims of major breaches, which led to millions of their users' information being compromised.

But breaches are also possible on a much smaller scale, for instance as a result of phishing emails. In a classic MO, a fraudster sends an email to an employee posing as a colleague, requesting data on other employees. It sounds simple, but it's startlingly effective; Snapchat fell victim to such a scam in 2016, resulting in compromised accounts of 700 employees. Another trick that fraudsters use is to call a merchant's customer service, and pretend to be a real customer (gaining trust by using simple PII they can easily obtain from social media). Eventually they try to convince the representative to hand them the account credentials.

Once login credentials are stolen, fraudsters need to test them to see if they're valid. Customers might have changed their logins recently, but fraudsters also want to check if the username/password combinations they have for one shop might work in others as well. Since people frequently use the same credentials for many different online accounts, this method often proves successful– and this tends to compound the damage of a breach exponentially.

But the most efficient way to test stolen credentials is by using bots. A bot is a simple software application that automates a task. Hackers use bots to 'credential phish' – test logins and passwords automatically and at extremely high speed until they successfully log in, thereby validating the credentials they've stolen.

Once a set of credentials is validated, a hacker can either use them themselves, or sell them on the Darkweb. The surprising affordability of this data on the open market is an unnerving testament to its abundance, as well as the efficiency of credential phishing: logins to Paypal accounts with a balance of $500 cost only $6.43, and Uber account logins cost under four dollars. And then the fraudster just has to decide how to best exploit their unauthorized account access.

What do fraudsters do once they're logged in?

It's frightening how many different avenues a fraudster can take to try to profit, once they have verified credentials. Some possibilities include Ransom Attacks or taking out loans in the user's name, using billing information and other PII

However, the above MOs are less relevant to eCommerce merchants, whose primary ATO threat is fraudsters logging in using real credentials and ordering goods. By the time the merchant and rightful account owner discover what happened, any information the customer has stored in the account is compromised, plus the merchant has shipped the stolen goods, and will be liable for a chargeback.

So how do fraudsters pull this off? Occasionally a fraudster will be lucky enough to have the credit card details of the same user he's logged in as. But Riskified has found that this occurs pretty rarely. The two most common patterns we see are loyalty fraud, and the fraud MO we call "Mismatched ATO".

“Mismatched ATO”

This is the most common pattern we see in ATO attacks. In these cases, a fraudster obtains account information, but not the associated credit card details. So to carry out the attack they use a stolen card card that belongs to an unrelated person.

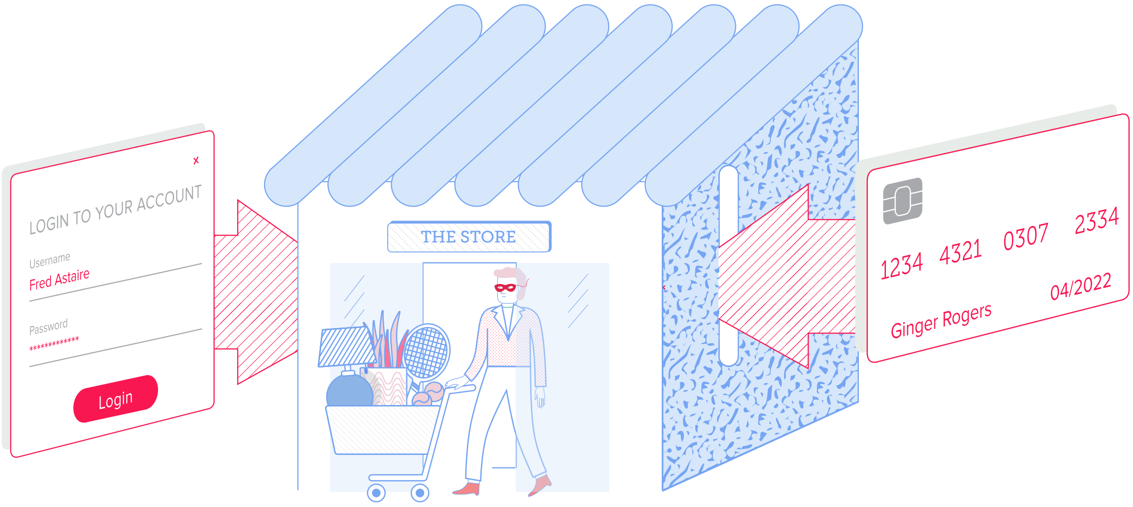

This sort of ‘Frankenfraud’ sounds like it should create enough of an incongruity to raise eyebrows, but in fact this attack traditionally has a high rate of success. Many merchants, unaware of the scope of the ATO issue, decide that good login credentials are enough to essentially auto-approve an order. And even when merchants detect something suspicious in one of these orders, they tend to refrain from requesting this customer take steps to verify their identity. If this really is the customer, merchants feel verification methods may tarnish the shopping experience and could send the shopper to the competition.

In essence, in a Mismatched ATO, fraudsters take advantage of merchants’ apprehension about reaching out to customers. More often than not, it works.

Loyalty fraud

As simple as the Mismatched ATO is, life can be even easier for a fraudster. If there is already some sort of store credit or rewards cash balance in the account, fraudsters can use it to shop immediately. The most common examples of this are frequent flyer miles or hotel loyalty programs, where it’s quite common that customers store significant amounts of value in their accounts.

When a fraudster commits loyalty fraud, the merchant is responsible for reimbursing that stolen store credit. Plus, the fallout is incredibly embarrassing for a merchant. Once a customer finds out that their account was compromised, a program designed to promote customer loyalty is likely to result in just the opposite.

Detecting and preventing ATO attacks

It's possible to implement ATO detection and prevention at various parts of the shopping process. Most merchants review orders for ATO as part of their general CNP fraud prevention process, which takes place at checkout.

But if ATO is detected at checkout it's too late, because the bad actor has already pilfered the customer's data. Even if you succeed in catching an ATO at the point of checkout, you're still obligated to inform the customer that their account has been compromised, so that they can change their login credentials, and make sure to cancel any payment options they've saved on your site.

The only way to truly protect your brand—and customers—from ATO attacks is to detect bad actors at the point of login. Admittedly, this makes things a bit trickier: First, a false decline at the point of login is even more insulting and confusing to a customer than at checkout, so precision is crucial. Second, your decision about how to respond to the login needs to be *instant* . While some customers may be willing to wait a minute or two while their order is reviewed, they won't put up with a delay while simply trying to access their account.

LOGIN TO YOUR ACCOUNT
Username
John Doe
Password
Login
PROCEED TO CHECKOUT
Credit Card No.
1234 1234 1234 1234
Exp. Date
11 / 04 /24
Purchase

The need for accuracy and real-time decisioning means there is no choice but to fully automate this process. Building such a system is not simple, but there are two main facets of any automated ATO detection system: detecting changes in behavior and device, and bot detection.

Behavioral & device change

The principles behind this method of spotting ATOs is similar to traditional fraud detection: gather information about the shopper and their device, and measure it against the database of historical, legitimate orders and logins you have for this customer. Fortunately, if someone is logging into their account, you've seen them at least once before (when they created the account).

Your automated detection system should note factors like:

- If the shopper logs in from a different geographical IP address than usual (or using a proxy server)
- If the shopper logs in using a different device than usual (you'll have to use device fingerprinting to figure this out)
- If the shopper logs on at different time of day than usual
- Password entry behavior (fraudsters tend to copy paste, legit customers type them in manually)
- Browsing time before login (fraudsters and bots login quickly, legit customers only do once it's required)

However, just because your system detects anomalous behavior doesn't mean this login attempt should be automatically blocked. Below, we'll discuss how to devise a verification strategy to deal with risky appearing login attempts.

Bot detection

In the case of credential phishing, where a bot is attempting logins at high velocity in order to test or guess credentials, it's very important to stop the bot in its tracks. If you let the bot complete the phishing process and verify login information, your business is obligated to inform the rightful account owner that their credentials have been compromised, which will make you appear incapable of providing adequate security measures.

An automated system can determine with fairly high accuracy if a shopper is a bot or not based on parameters like:

- Keystroke velocity
- Velocity of login attempts
- Mouse movements
- The way the user scrolls
- Mobile device orientation sensors

If your system determines that it is a bot–not a human–attempting to log into your store, then you know it's likely a case of credential phishing. Note, however, there are certain cases where "real" customers use bots, and you may want to allow them to log in and check out. A good example is releases for new lines of sneakers, or limited edition cosmetics. Savvy shoppers will sometimes employ bots in order to checkout faster than a human possibly could, ensuring that they secure the products during this limited release. Whether or not to allow these bots into your store is a policy decision, not an issue of fraud detection.

Catching 100% of bots and fraudsters is a tall task, and it's likely that your detection accuracy won't be perfect, given the pressure of time constraints. For this reason, it's important to conduct retrospective analysis on chargebacked orders to understand if you were a victim of an ATO attack. It's not as good as catching a fraudster at the first attempt, but at least you can ensure that you don't fall victim to the same fraudster again. Furthermore you can inform your customers that their credentials have potentially been compromised in order to curb further damage.

Post detection: How to respond

Even the best ATO detection system will rarely reach a 100% clear cut decision on whether a login attempt is legitimate or not. Think of it like waking up in the morning and trying to determine whether it will rain or not based only on certain data points: location, time of year, the way the sky looks. Some mornings will be fairly clear cut, others will be edge cases, where you're pretty unsure.

The downside of guessing wrong is pretty severe. Allowing a fraudster to log in means compromising your customer's data and exposing yourself to financial loss. But being too cautious and blocking a good customer's login attempt is likely to alienate them forever.

So in a sense, detection is only half of the equation. The next step is determining thresholds of certainty for outright blocking vs. further identity verification vs. allowing the customer to log in. Yes, verification measures risk creating friction in the buying process, and can lead to dropoff. But when it comes to ATO detection, selected verification measures are necessary for edge cases, where your system simply doesn't have enough information to make a confident decision. In fact, in cases where customers have entrusted you with a lot of sensitive data or store credit, some verification may reassure them that their information is safe.

First you need to decide on a threshold of risk. How certain must you be that a login is legitimate to allow it without verification measures? The answer will depend largely on what kind of data you store for your customers. If you're an airline and your clients store valuable frequent flyer miles in their accounts, or your customers keep their entire credit card numbers stored for easy payment, it pays to be risk averse.

Once you've answered these questions, it gets more complex, because not all verification measures are the same. Once an automated solution decides that data around a specific login attempt is inconclusive and merits a follow-up, it then needs to decide from a range of possibilities including: sms messages, Captcha, emails, email-based login alerts and security questions. Again, the precise mapping of a verification measure to a risk profile depends on an individual company's situation, but there are some considerations merchants should definitely be aware of.

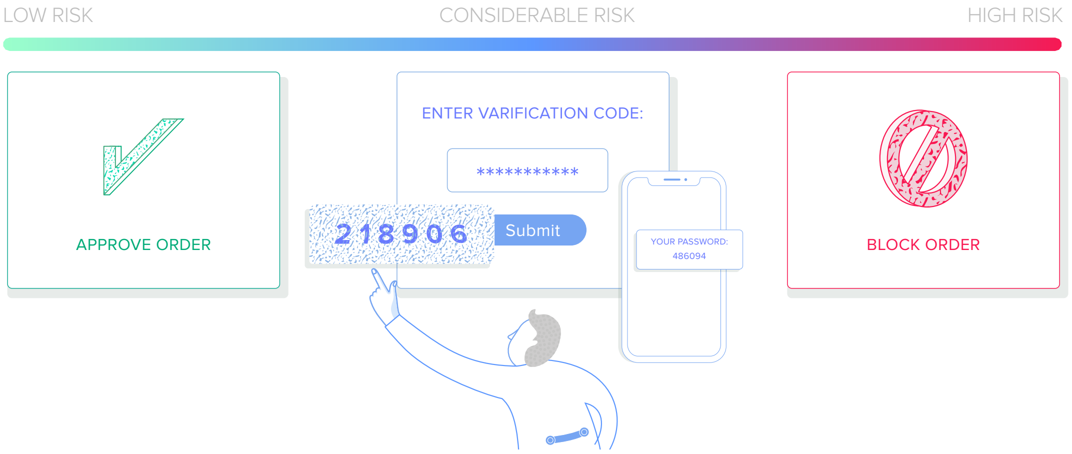

SMS messages

Obviously, for an SMS to be an effective way of verifying identity, you need to make sure you're sending it to the phone number of the real account holder. Ideally, merchants should make the phone number a required field during account creation for this purpose. The second best option is to require customers to enter their number when making a purchase. Then you can confirm that the number they entered is the right one by cross-checking it against the information entered in a previous (legitimate) transaction. It's also best practice to search for a match between the phone number and the billing address you have on file, using a service like Whitepages.

Then there's the decision about the content of the message. We recommend sending a verification code, which the customer needs to enter on your eCommerce site in order to login. This is better than a message which requires them to reply to a text confirming they are the shopper, because the former forces the customer to re-engage with the site, meaning it's less likely to lead to drop-off.

Email verifications and alerts

Like with text messages, an email verification is only worthwhile if you're sure the email is being delivered to the account holder. You should ideally send an email to the address used to open the account. However, because this is an ATO attack, you need to consider the possibility that the email has been hijacked.

To verify that the email is still in the user's hands, we recommend sending a message alerting them that there's been a suspicious login attempt to their account and requiring the recipient to take some action like clicking whether they recognize this attempt (providing yes/ no buttons). Once someone clicks a link in your email, you can compare the IP address to the location provided when the account was first opened.

If you're more worried about customer friction, or less worried about this particular login being an ATO, you can just send an alert email which requires no action, to inform customers that there was a login attempt from an unrecognized device. You may notice that Google does this when you buy a new phone and open your gmail the first time.

Conclusion

Because ATO attacks impact returning customers, they pose a unique threat to merchants' eCommerce revenue. Failing to prevent ATO fraud can lead not only to chargebacks, but also puts your customers' privacy and your brand reputation at risk. Unfortunately, ATO attacks are likely to become increasingly frequent and sophisticated as cybercriminals adapt their activity to try and outsmart fraud prevention systems.

Collecting and analyzing data about users' online behavior in real-time and comparing it to customers' past behavior is the key to detecting ATO attempts. But identifying the attacks isn't enough. To successfully manage this type of fraud, merchants need to determine when and how to block bad users, notify customers of suspicious login attempts, or request additional verification.

About Riskified's ATO solution

Riskified built an ATO prevention solution that accurately identifies fraud attempts, delivers actionable next steps, and minimizes friction for your customers. We use machine learning models, digital fingerprinting and state-of-the-art bot detection to assess the risk of every login attempt and provide a clear “allow”, “block”, or “verify” decision.

By linking every login attempt to millions of previous shoppers and billions of transactions across Riskified's merchant network, our system recognizes legitimate customers and bad actors with a high level of confidence, reducing the need for additional verification. In the case that a login attempt—or other suspicious behavior during the shopping journey—requires a closer look, merchants can choose to have Riskified automatically deploy the appropriate verification measure.

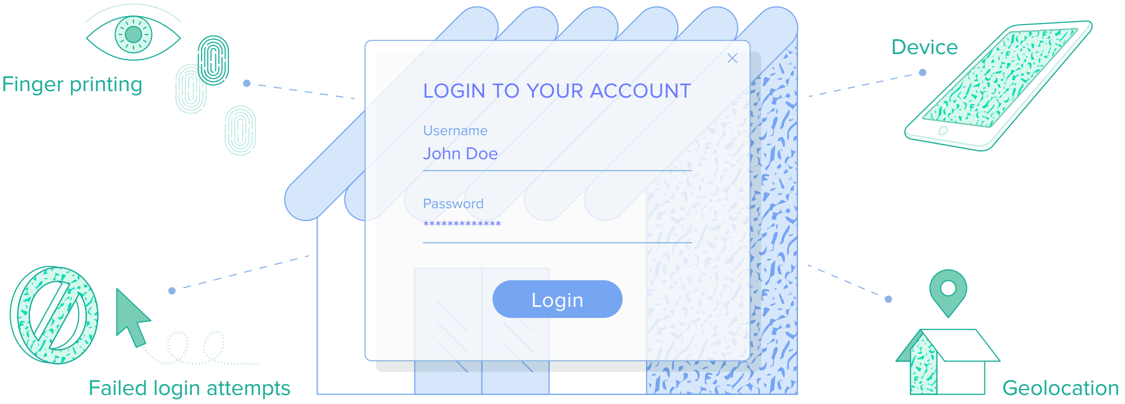

Merchants who partner with Riskified benefit from best-in-class fraud detection technology, actionable decisions, and ensure that their good customers are never subjected to needless friction.

For more about Riskified's ATO prevention solution, contact sales@riskified.com.