neuroID

A part of experian

Report

Volume 2: The power of one

THE IMPOSTER IN THE SESSION

An investigative look at the fraud rings abducting your customers' accounts.

Table of contents

Fraud rings are typically painted as large, coordinated groups; their attacks are described as massive, coordinated invasions that descend on a single target. The assumption is that the damage they inflict is directly tied to their size — the larger the ring, the bigger the attack they can launch.

The threat of fraud rings rightfully draws fraud leaders' attention. 71% of businesses identified fraud rings or financial criminals as the source of attempted fraud in 2025,[1] with account takeover (ATO) attacks in particular growing 24% year-over-year.[2] Almost 3/4s of businesses are planning to boost their fraud prevention budgets in response to rising losses[3] — based on the attention fraud rings garner, stopping coordinated attacks is likely to be a major area of investment.

But the staggering size of the world's largest fraud operations often overshadows reality. **The impact of fraud rings doesn't always manifest as large-scale attacks, and coordinated ATOs aren't the only way for fraudsters to capitalize on evolving ATO tools.** Fraud rings are often behind the scenes of seemingly isolated attacks — supplying stolen credentials, selling attack kits and publishing playbooks that make solo operations possible.

In this report, we'll zero in on another ATO attack in action. This one looks distinctly different from the coordinated hub-and-hive structure we examined in volume 1, but its impact shows why the danger of fraud rings extends beyond large-scale invasions.

How NeuroID spots fraud rings in action

NeuroID uncovers ATOs using four layers: **persistent device recognition, geolocation, behavioral intelligence and multidimensional link analysis**. We recognize devices we've seen before, with 99.5% accuracy, then check location signals and behavioral patterns. If a device appears in a different location from the account or the user's behavior is identical across multiple accounts (copying and pasting credentials, navigating the profile page), it signals coordinated fraud.

When risk spikes, **we collaborate with our customers to investigate**. We map connected devices and accounts first, then layer in masking signals like VPNs, proxies, Apple relays, TOR nodes and GPS spoofing, and finally analyze intent: Are users navigating with precision, copying credentials or repeating steps across accounts? We pull these layers into network graphs to reveal the size and sophistication of the attack. While fraud rings continue to spoof networks and devices, behavior remains the hardest to fake at scale — making it the ultimate safety net against ATO rings like the one featured in this report.

Family, friends or fraud?

Let’s start with two devices logging in to a single account at a payment provider. Joint tenancy, or account sharing, is common in this industry, so it’s not unusual for accounts to have multiple devices connected to them.

Separating a legitimately shared account from ATO can be extremely difficult without the right tools. The following attack takes advantage of that, defying what we’d expect from a coordinated attack to fly under the provider’s radar.

Account 1 is accessed by two devices — **device A and device B** — during the same 30-day billing period.

The devices are located in different areas (device A in Indianapolis, Ind. and device B in Charlotte, N.C.) but take the same actions: They both log in with a password and add a new payment card to the account. Besides location, the biggest difference between the two is that device B logs in four times during the billing cycle, while device A only logs in once.

In a vacuum, this looks a lot like joint tenancy. Maybe this is a couple treating account 1 as a joint account, or a parent adding their payment card to a child’s account. The location difference could easily be due to normal travel. Without diving deeper, there’s no reason to suspect high risk and therefore place excessive friction in front of either of these users.

But there are red flags that lie beyond these logins. These devices’ activity reveals one as the likely true account owner, and the other as an imposter masquerading as a genuine user.

Account 1 is accessed by two devices: device A and device B. In a vaccum, there’s no reason to place excessive friction in front of either user.

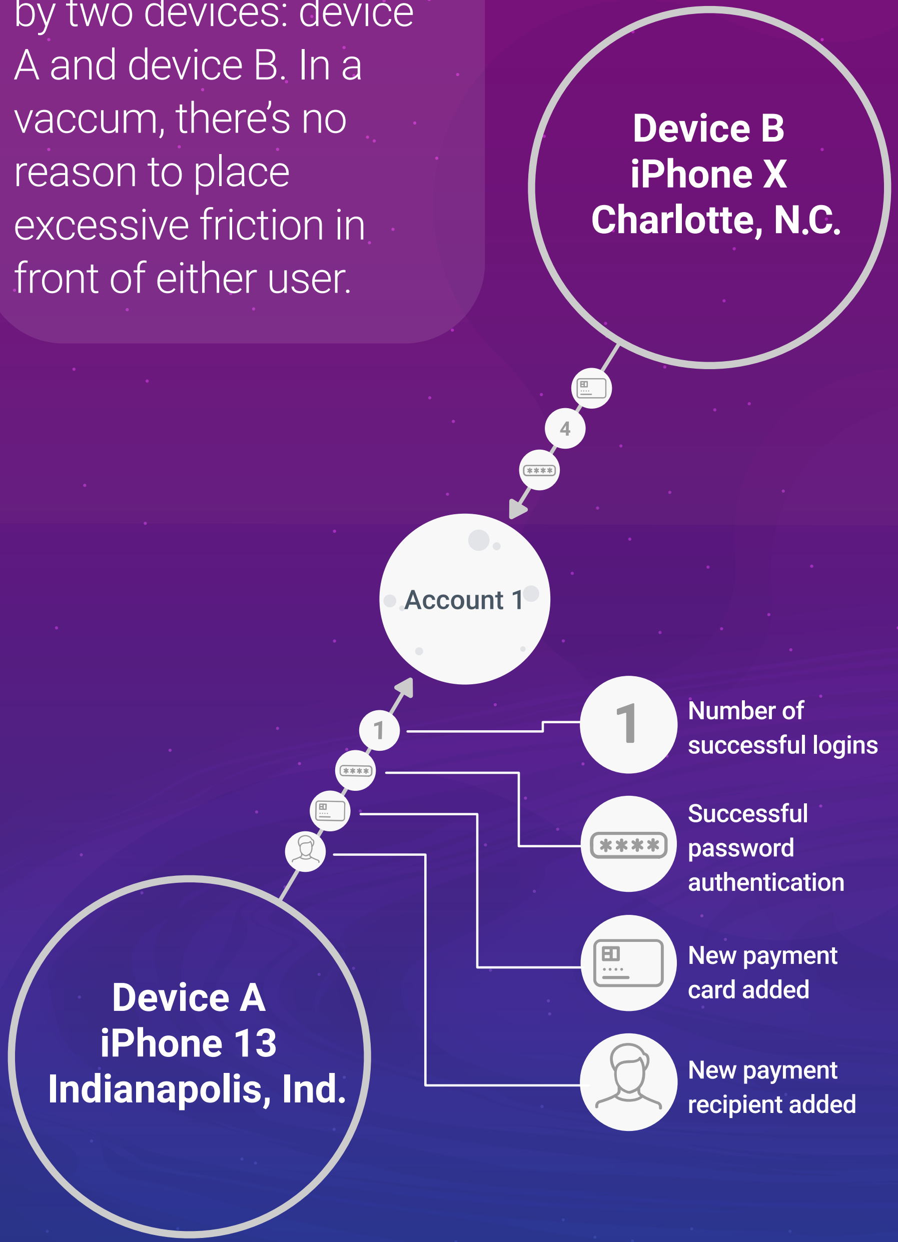

A high-speed takeover

As we uncover more about these devices, the imposter quickly reveals itself.

NeuroID traced each device's activity, revealing a clear dichotomy between device A and device B. **Device B** only logs in to account 1 and does so multiple times — a low-risk profile that indicates device B is account 1's true owner.

Device A, on the other hand, accesses 50 additional accounts within this billing cycle. NeuroID recognized device A on all 51 accounts. As we studied device A's actions across all the accounts that it accessed, a pattern emerged — device A took the same actions on nearly every account:

Step 1: Made a single successful login using a password.

Step 2: Add a payment card.

Step 3: Add a payment recipient.

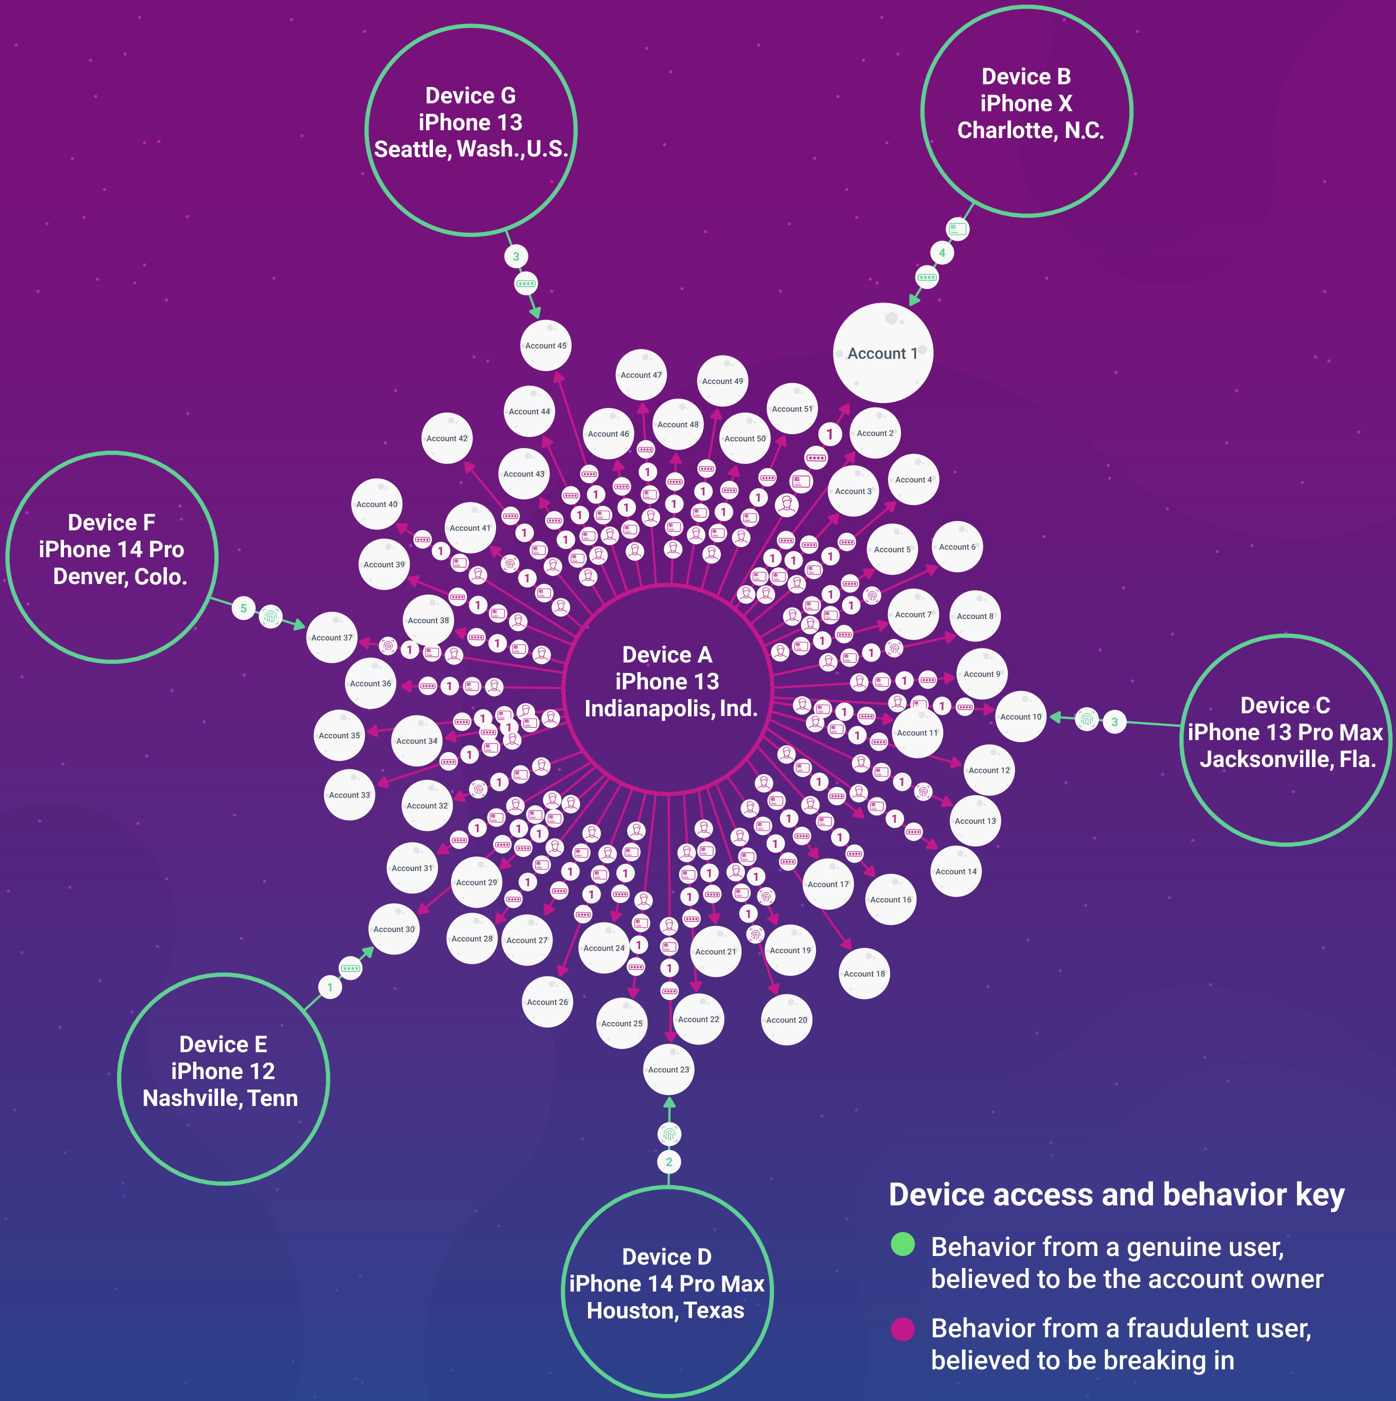

Device access and behavior key

- ● Behavior from a genuine user, believed to be the account owner
- ● Behavior from a fraudulent user, believed to be breaking in

In the payments space, these actions are common in fraud attacks: Fraudsters access accounts, then add and use stolen payment cards to funnel funds to themselves or their associates. The taken-over accounts serve as a vehicle for stealing or laundering funds.

In this specific attack, these procedural steps are noticeably quick and low-touch. The fraudster(s) behind device A don't try to build history between the accounts and their device. Rather, they're working quickly to monetize as many successful ATOs as possible.

There are a few accounts where, like account 1, the genuine user logs in multiple times during the billing cycle.* In each case, the true owner's device is located in a different part of the U.S. than device A, but device A doesn't use a VPN to mask their location — a low tech, common tactic fraudsters use to appear more like legitimate account owners. Instead, device A continues to prioritize efficiency, disregarding even minimal masking: Setting up and changing a VPN for each login attempt takes time, which this fraudster shows no interest in wasting. This strategy is called a brute force or blitz.

*The other accounts have true owners, but weren't logged in to by them in this window.

A one-man ring?

This attack looks different than the coordinated fraud ring we last examined. Device A seems to be working alone; there are no "branches" connecting it to other fraudsters.

The structure of the attack is smaller, but the steps taken to execute it are largely the same. Device A is confirming credentials by logging in, then monetizing the takeover by adding payment cards and recipients.

Even though we don't see a network of fraudsters in this attack, that doesn't mean it's not a fraud ring's work in action. Fraud rings create and sell playbooks that teach individuals how to replicate ring tactics without the overhead of coordination. When device A takes the same, repeated actions on each account, it's likely following a step-by-step walkthrough to execute the attack. These playbooks are tailored to maximize a single user's efforts, **putting the power of a coordinated fraud ring in the hands of an individual**.

Behind the lone wolf

How fraud rings power solo attackers

Through the growing fraud-as-a-service industry, fraud rings empower individuals to replicate their tactics in a number of ways.

Playbooks for purchase: Detailed guides outline every step of an ATO attack, from credential testing to monetization. These playbooks are sold on dark web marketplaces and optimized for individual, speed-first attacks.

Stolen credentials: Rings supply the raw materials for solo fraud — batches of compromised usernames, passwords and payment card data harvested from breaches or phishing campaigns.

Attack toolkits: Beyond credentials, rings sell automation scripts and credential-stuffing tools that let individuals scale attacks without technical expertise.

VPN and spoofing instructions: Playbooks often include guidance for masking location and device signals, helping lone actors mimic legitimate users and evade detection.

Community support: Fraud forums and encrypted channels allow individuals to share tips, troubleshoot issues and refine tactics — creating a distributed network of "solo" attackers who are anything but isolated.

The strength of the solo fraudster

The key difference between a solo attacker and a coordinated ring? Quality versus quantity. A ring with multiple members and assigned roles can spend more time testing defenses, leveraging tools to bypass them and cementing control over accounts. Their size allows them to scale attacks quickly without compromising effectiveness.

An individual fraudster doesn't have that ability. If they adjusted their tools and logged in to each account dozens of times like a coordinated fraud ring, their impact would be severely limited by time and resources.nstead, individual fraudsters take a lower-touch approach, meaning — if you're looking for the signals — they more clearly stand out when their device data, location and behavior are compared to genuine account owners. As a result, step-up authentication is more likely to stop an individual fraudster than a coordinated ring member.

But the speed of solo attackers is also their biggest advantage — and the playbooks created by fraud rings and sold to individual fraudsters are designed to maximize it. By opting to target more accounts (with less time invested in each), individual fraudsters move faster and maximize the limited resources they have. It's an intentional effort that trades total control over a few accounts for a foot in the door in many. The sheer number of accounts accessed means that the individual fraudster has the potential to do as much damage as a midsized ring, even if step-ups limit their success rate.

Take the last attack we studied: that attack had four devices collaborating, taking their time to cement control over 20+ accounts. The individual in the attack above accesses over 50 accounts. Even if they're stopped by step-up authentication on half of their attempts, **this individual would inflict $7,500 in losses* — the same as the coordinated fraud ring**.

*Based on the provider's average loss of $300 per successful ATO

The fraudsters' marketplace

On the dark web, individual fraudsters can buy data tools to execute ATO attacks — the same resources used by major fraud rings.

Attack playbook: **$50**[4]

Verified credentials: **$30**[5]

OTP relay: **$15 per attempt**

Credit card with CVV: **$15**

Stopping the individual imposter

Identifying the individual imposter is just as important as stopping a coordinated ring's invasion. The lone fraudster uses its individuality to its advantage, scaling their efforts through a fast-paced approach while remaining detached from a larger operation that could give them away. This is a more extreme case, but imagine one where a fraudster only logs in to three or four accounts. How would you reliably separate it from a real user accessing shared accounts?

The best way to weed out the individual imposters is to identify high-risk actions like these: repeated steps to set up fraudulent transactions, location mismatches with other account users, total number of accounts accessed and more. Even when there's not a larger fraud ring connected to it, stopping these solo devices early prevents attacks from scaling.

To put it in practice in this scenario: There's no reasonable explanation for a genuine user to access 51 accounts, so that alone would be a reason to stop device A. But even if its login to account 1 was the first time we saw device A, there are red flags: device, location and behavioral mismatches, plus known fraudster tactics for monetizing ATOs. Putting a barrier in front of device A — flagging it as risky and requiring a step-up authentication for subsequent login and transaction attempts or outright blocklisting it — would prevent the device from accessing the other 50 accounts.

Flagging device A as risky on the first login attempt — therefore requiring step-up authentication for subsequent login and transaction attempts — would prevent the device from accessing the other 50 accounts.

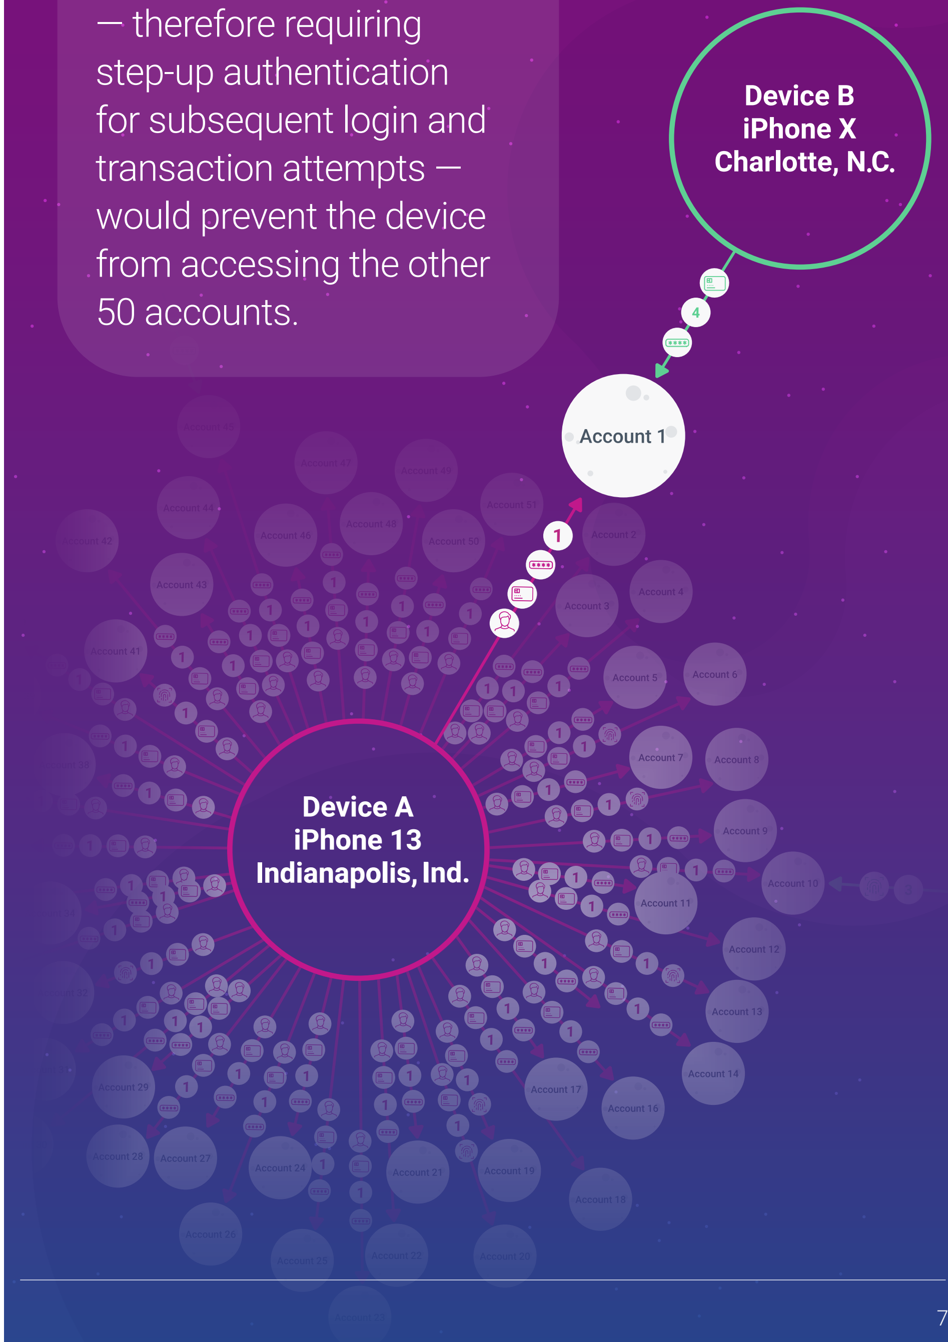

Shutting down attacks of all sizes

This attack is another example where multifactor authentication (like one-time passcodes sent via SMS), biometrics (including fingerprint and face scans) and similar point solutions fall short by treating risky logins as individual events. These speed-first solo fraudsters know they'll face friction on some attempts, but even a modest success rate results in a major attack.

NeuroID stops both individual fraudsters and coordinated rings by looking at attacks differently. Through triangulation of behavior, device intelligence and geolocation, overlaid with multidimensional link analysis, NeuroID exposes and stops ATO attacks of all kinds. Whether it's a lone wolf or a large pack, we shut down the infrastructure that makes strategic, coordinated ATO attacks possible.

Want to see more fraud ring playbooks in action?
Read the first installment in this series, and visit neuroid.com to learn more about NeuroID's ATO protection.

1. Alloy 2025 State of Fraud Report

2. Beyond the Breach: Account Takeover Data & Insights, Sift, 2024

3. 2025 U.S. Identity & Fraud Report, Experian, 2025

4. What We Learned From Infiltrating 22 Credential Stuffing Crews, Kasada, 2025

5. Dark Web Price Index, Privacy Affairs, 2023