neuroID

A part of experian.

# The Fraud Attack Strategy Guide

How to read fraudsters' battle blueprints and use their tactics against them

"Let your plans be as dark and impenetrable as the night, and when you move, fall like a thunderbolt."

- The Art of War, Sun Tzu, 500 BCE

"Know your enemy and know yourself, and you need not fear the result of a hundred battles."

– Sun Tzu

You're at the forefront of the fight against fraud. Fraudsters deploy coordinated attack strategies, test defenses, and adapt to countermeasures, while fraud prevention teams must anticipate attacks and stay ahead of evolving fraudster tools and techniques.

This back-and-forth isn't new, but the Generative AI (GenAI) arms race has changed the rules of engagement.[1] At the beginning of 2024, fraud bots comprised 30% of all fraud attacks observed by NeuroID; **by the end of 2024, bots were 80% of all attempts NeuroID observed, nearly tripling in frequency in less than a year** according to a NeuroID internal analysis. GenAI is creating bots that are scalable and aggressive beyond any previous iteration. And that's only the beginning. The terrain and tactics of the digital battlefield are evolving faster than ever.

## Reading the Fraudster's Battle Blueprints

The rise of automation, artificial intelligence, and decentralized networks has accelerated fraud armaments. Attackers work in well-organized networks with specialized roles, defined objectives, and coordinated attacks.

In *The Art of War*, the famous treatise from the 5th century BCE, author (and renowned general, strategist, and philosopher) Sun Tzu emphasizes that victory belongs to those who understand their own strengths and weaknesses as well as their opponents.[2] **This principle is as relevant to your fight against fraud as it was to Sun Tzu's battle plans 2,500 years ago.**

Sun Tzu also teaches that the best way to win a fight is to avoid unnecessary battles — to prevent conflict by making attacks too costly or difficult for the opponents. In the digital world, this means creating multilayered stacks where fraud is not just detected but used as a lesson towards futureproofing against new threats. By treating each attack as a glimpse into the fraudsters' greater strategy, you can anticipate and undermine new fraud tactics, no matter how complex.

1. https://www.neuro-id.com/resource/blog/who-is-winning-in-the-race-for-genai-weaponization-fraudsters-or-fraud-teams/

2. Sun Tzu, The Art of War.

**" Whoever is first in the field and awaits the coming of the enemy, will be fresh for the fight; whoever is second in the field and has to hasten to battle will arrive exhausted. "**

– Sun Tzu

The fraud battleground is vast, and the opponents are complex. But NeuroID is a veteran of the field, and our behavior and device solutions were built to adapt faster than fraudsters. By analyzing our fraud attack data — the timing, causes, effectiveness of response strategies, and more — we can break down the primary attack blueprints and in-the-trenches strategies of today's fraudster. This research provides a view of your enemies' battle plans, including the new levels of exploitation, trust manipulation, and technology weaponization behind every third-party fraud attack.

**Understanding these four fraudster attack blueprints will help you turn the tide of battle in your favor.**

# Contents


Attack Style: Reconnaissance-in-Force 4

Attack Style: The Blitz 7

Attack Style: The Feint 10

Attack Style: Adaptive Tactics 13

# Attack Style: Reconnaissance-in-Force

In this attack, fraudsters test their enemies (you). They purposefully provoke you into revealing vulnerabilities, then exploit those vulnerabilities with a massive strike. **Here's how you stop them.**

> **"If your opponent is of choleric temper, seek to irritate him. Pretend to be weak, that he may grow arrogant."**
>
> – Sun Tzu

## Story from the Frontline: Fraudsters Probe for Weakness Before Striking

In the summer of 2024, a major credit card issuer experienced an unusual spike in application traffic. Given the traffic spikes in the flanking weeks, our customer wasn't surprised by the spike and had expected to see some fraud. However, NeuroID noticed the later spikes had a higher composition of risky traffic, meaning the issuer experienced a fraud attack that was hidden among the spiking genuine traffic.

NeuroID dashboards showed the volume spike was disproportionately comprised of risky applicants, continuing at anomalous levels for weeks. This was in fact a coordinated reconnaissance mission by fraudsters.

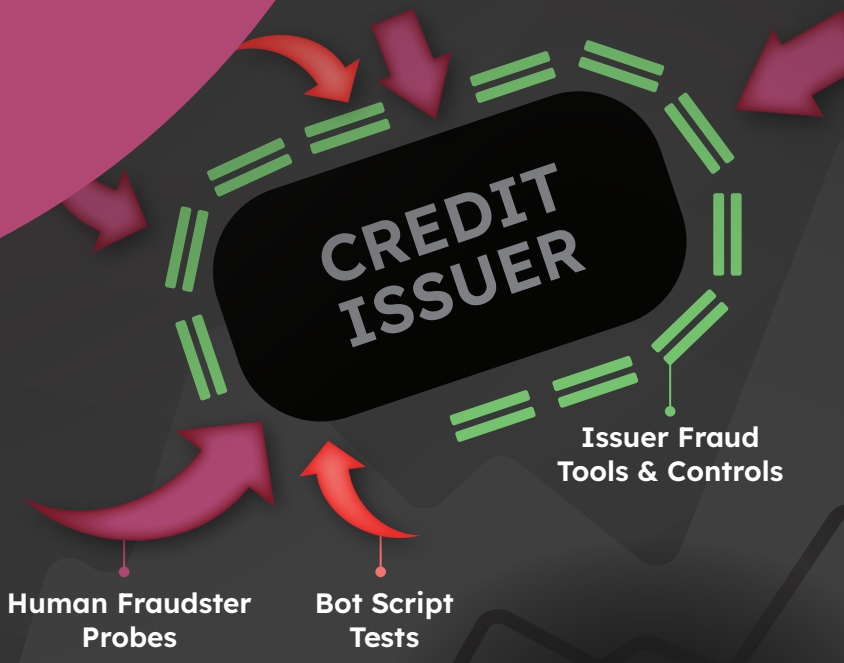




neuroID
A part of experian

## Tactical Breakdown

Among the decreasing spikes in application volume (visual 1), one spike was actually 15% bots. But this wasn't a single-day blitz attack. It was a carefully executed reconnaissance-in-force.

Starting in early August, we saw the fraudster traffic grow — starting with some smaller, discrete bot attacks followed by the nearly 15% of total volume one day mid-August (visual 2). This was probing, a strategic build up of tests ahead of the focused attack day.

For this issuer, we saw multiple deployments of bot groups over multiple days, leading to a total increase of 11K bots across August — well above their 2K baseline (visual 3). Bot script segment counts changed during each deployment as human fraudsters adjusted their scripts. **Their objective was clear: identify fraud controls, weaknesses, and ideal paths for future attacks.**

1 | When looking across all traffic, can you identify which spike(s) was almost 15% bots?

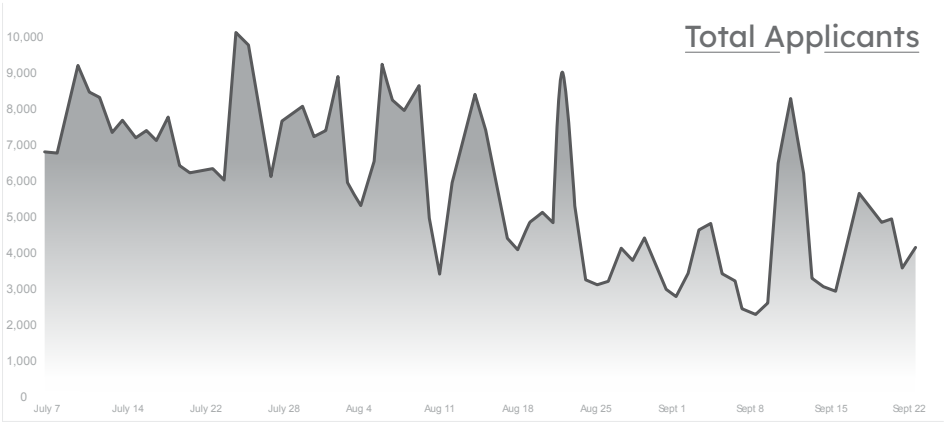


2 | It's here, sneakily growing after a spike of genuine applicants (green) starting on July 22. The traffic spike should have gone down after a week, but with the increasing fraudster activity (red), it stays for nearly a month.

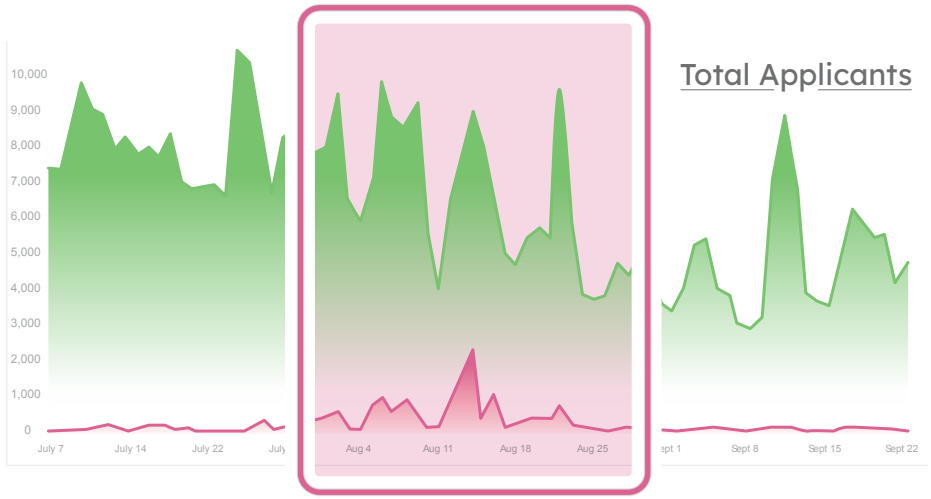


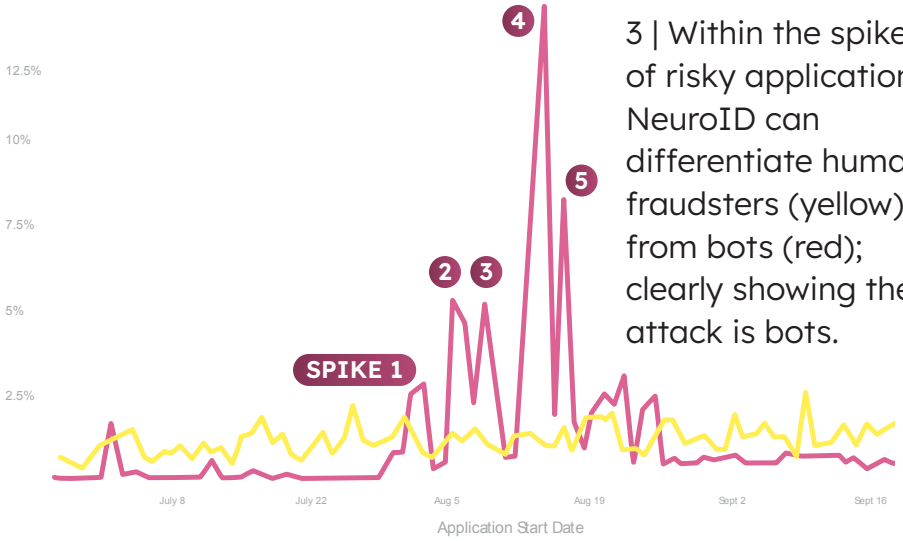


3 | Within the spike of risky applications, NeuroID can differentiate human fraudsters (yellow) from bots (red); clearly showing the attack is bots.



neuroID
A part of Experian

The first reconnaissance mission succeeded (see spikes 1, 2, and 3 in visual 3). With that intelligence in their back pocket, fraudsters launched a large-scale attack days later (see spikes four and five). Applications associated with the large-scale attack had a 70% application submission rate — higher than that of genuine users. These 11K bots should have been kicked out of the application pre-submit to reduce the financial cost and resource burden of manual reviews.

## Countermeasures

The first countermeasure is to identify the device and network IDs associated with the probes to prevent further testing — often probes have the same device or network IDs, making it easy to stop a cluster of 5-10 if you can identify at least one. Next, study the areas in the application the probes focused on, particularly drop-off areas; there could be a vulnerability exposed by the probes that needs to be repaired ASAP. Finally, ensure your fraud team is on high alert for a swell of unwanted attackers.

This credit card issuer was trialing NeuroID capabilities when the fraudsters' reconnaissance attack began. With our behavioral visibility, their fraud prevention team was able to begin detecting irregularities **five days before the full attack.**

## Actionable Intelligence

Without behavioral analytics, these reconnaissance missions testing for vulnerabilities would go under the radar at worst (or at best, would create a backlog of unnecessary manual reviews). With next-gen bots capable of unprecedented large-scale attacks[3] and Synthetic Identity Fraud (SIF) now making up 85% of all fraud, attacks like this are going to be even more frequent.[4]

With behavioral analytics, you can turn fraudsters' probe attacks against them and use these spikes as a warning sign. They're the slight troop movements over the horizon that you can see only if you've got the high-tech night gear. By detecting these probes in real-time, you stop not only the coming attack but deflect others that could potentially exploit the same vulnerability.

3. https://www.neuro-id.com/resource/new-industry-report-are-fraud-bots-beating-behavioral-analytics

4. https://www.neuro-id.com/resource/blog/5-fraud-forecasts-what-we-got-right-and-wrong-in-2024-so-far/

# Attack Style: The Blitz

In this attack, fraudsters launch rapid, concentrated assaults to overwhelm the opponent (you) before you can effectively respond. **Here's how you stop them.**

## Story from the Frontline: A Top Five Bank Under Siege

In October 2024, NeuroID dashboards detected one of the most aggressive fraud campaigns we've ever seen. It was targeting a top five domestic bank and was a classic blitz: flooding the system with an overwhelming scale of attackers before defenses could adjust.

> **"Speed is the essence of war. Take advantage of the enemy's unpreparedness; travel by unexpected routes and strike him where he has taken no precautions."**
>
> – Sun Tzu

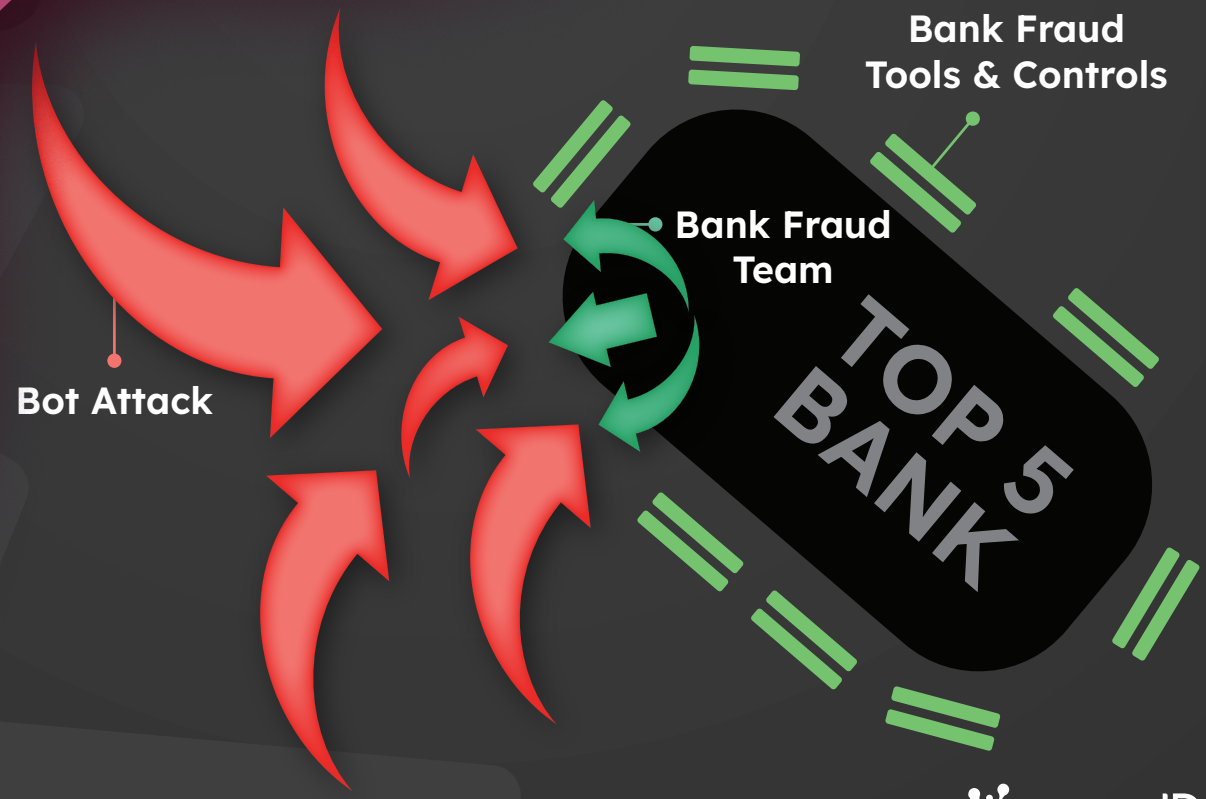




neuroID
A part of Experian

## Tactical Breakdown

On a single day in mid-October, bot traffic made up nearly 20% of total applications (visual 2). Like any blitz attack, the sheer volume was designed to overwhelm fraud mitigation forces in the hopes that even a small percentage of fraudulent applications would be approved. But even the most sudden attacks are preceded by scouts gathering intel before they strike.

Six weeks before the attack, we saw an uptick in human fraudsters followed by an increase in bots a couple of days later (visual 3; spike 1). This pattern repeated itself: human fraudsters and bots traded off repeatedly for a week leading up to the blitz attack (spikes 2 and 3). These human-to-bot transitions are a common theme we see: fraudsters are fine-tuning their scripts in preparation for the brute force attack.

The day after the blitz attack, we saw a significant decrease — but not total disappearance — of bot traffic. It even crept back up the week after the blitz, where there's another uptick in bot volume (spike 5; almost 8% of total volume).

1 | When looking across all traffic, can you identify which spike(s) were almost 20% fraud?

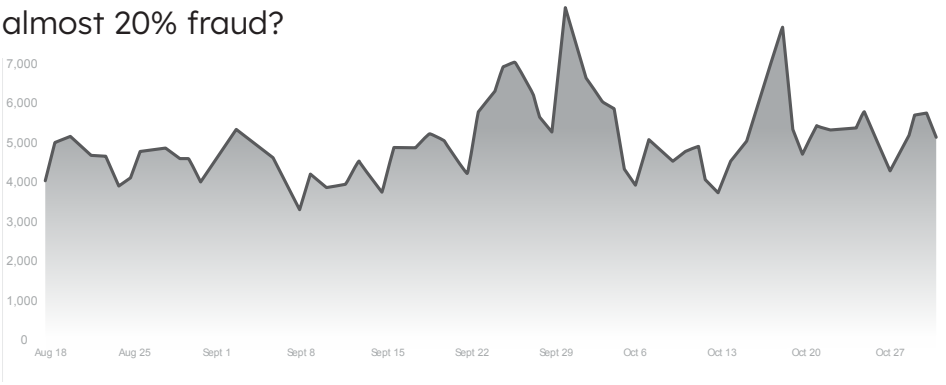


2 | It's here, generating a spike in daily application traffic, but not one higher than a couple weeks prior. Would this traffic spike alone have alarmed you?

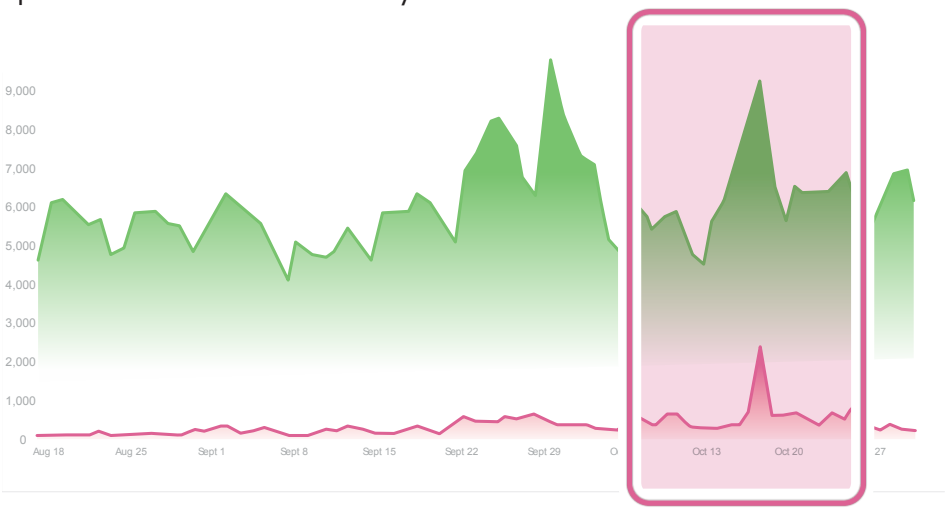


3 | See the small spikes of human fraudster activity preceding the bot spikes? These human-to-bot transitions show preparation and testing ahead of the mid-October attack.

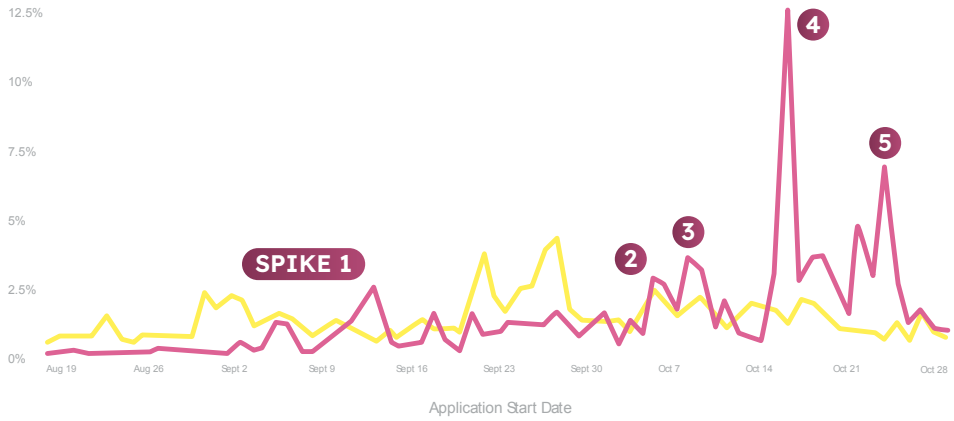




neuroID
A part of Experian

This could mean two things: either the blitz was successful enough to try again, or fraudsters assumed that the bank's fraud team was so overwhelmed with manual reviews that they could send in another attack without being noticed. This is another common fraudster tendency: they want to exploit known weaknesses for as long as possible and maximize their impact.

## Countermeasures

Blitzes often arrive with little warning, giving attacking fraudsters the high ground on defensive mitigation teams who aren't prepared for their speed and scale. Stopping them requires a fortified first line of defense. Effective top-of-funnel fraud detection that can auto-decline bots and human fraudsters without manual reviews is a necessity to survive today's fraud battleground.

We've seen customers hesitate to auto-decline risky applicants to appease customer acquisition and marketing teams, and it hasn't ended well; one attack victim had to shut down their entire online application while they worked through a backlog of manual reviews (read how we helped optimize their processes here).

In the bank's case, their fraud team turned to NeuroID to detect bots that their other defenses couldn't, revealing over 20K bots that no other tool in their fraud stack detected. They continue to use NeuroID to block bots at the top of their onboarding funnel and keep blitzes at bay.

## Actionable Intelligence

Speed, scale, and shock are key weapons in the fraudsters' arsenal. NeuroID behavioral analytics, combined with device and network intelligence, are the a proven way to detect these advanced blitz attacks and counterstrike in real-time.

NeuroID's countertactic to today's next-generation bots is advanced behavioral analysis, which catches bots more effectively than traditional device-based detection. As GenAI-powered bots become more sophisticated and harder to detect through conventional device or network signals, behavioral insights give you a critical advantage. NeuroID's advanced bot signal, specifically designed to identify next-gen fraud bots, has demonstrated a significant lift in bot detection for our customers[5] — especially in catching bots that device and network solutions alone would have missed. **If you still primarily rely on device and network tools to detect bots, you're trusting medieval defenses to stop your opponent's much more modern arsenal.**

5. https://www.neuro-id.com/resource/new-industry-report-are-fraud-bots-beating-behavioral-analytics

# Attack Style: The Feint

In this attack, fraudsters pretend to attack one vulnerability and draw the opponent's (your) attention and resources. Meanwhile, the real attack is happening where you least expect it.
**Here's how you stop them.**

## Story from the Frontline: Deception at Scale

A leading lending platform faced what they thought was a large-scale, albeit run-of-the-mill, bot attack in mid-2024. But the obvious threat was a misdirection hiding a far more insidious and difficult attack to detect.

> **"When we are able to attack, we must seem unable; when using our forces, we must appear inactive; when we are near, we must make the enemy believe we are far away; when far away, we must make him believe we are near."**
>
> – Sun Tzu

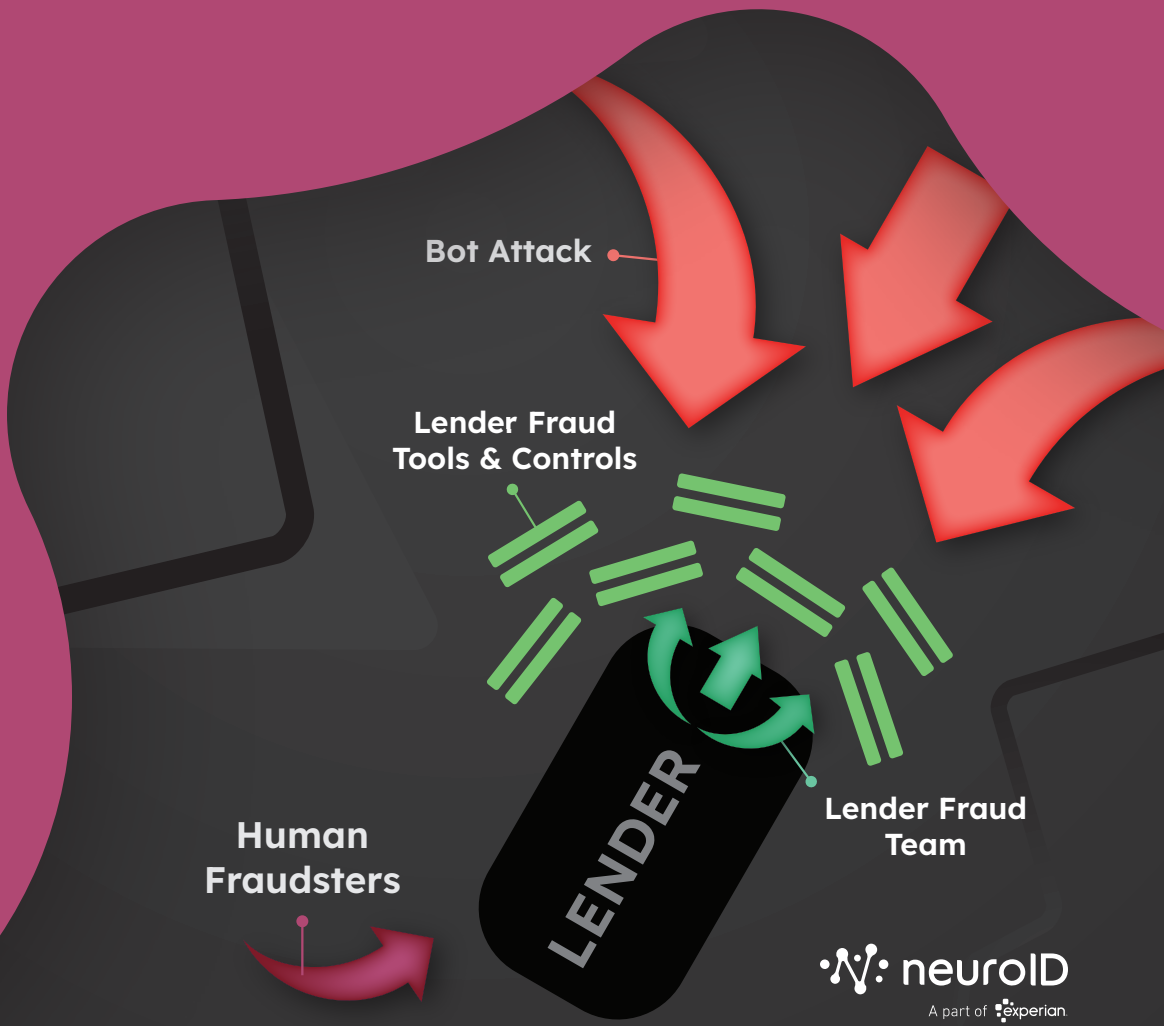

## Tactical Breakdown

This lender typically had very low bot activity, especially compared to their peers. So understandably, when bots started to make up more than 40% of their applicant pool, they were distracted. But NeuroID doesn't just see bots, we also see human fraudsters. And the human fraudsters started infiltrating their application a little bit each day, trying to avoid attention. At its worst, these human fraudsters ranged from 5 to 15% of daily application volume.

Despite their creeping success, they initiated a feint — a bot blitz comprising 40% of total application volume in one day (visual 2). This obvious attack diverted the attention of the lender's fraud team. While they investigated the bots, the human fraudsters went unnoticed and continued their covert operations. Why a feint? They could have been testing for vulnerabilities, trying to hide under the radar, or concerned they'd be found out soon.

1 | There's an obvious spike in fraudulent activity. But when looking across all traffic, can you identify when the fraud attack truly began?

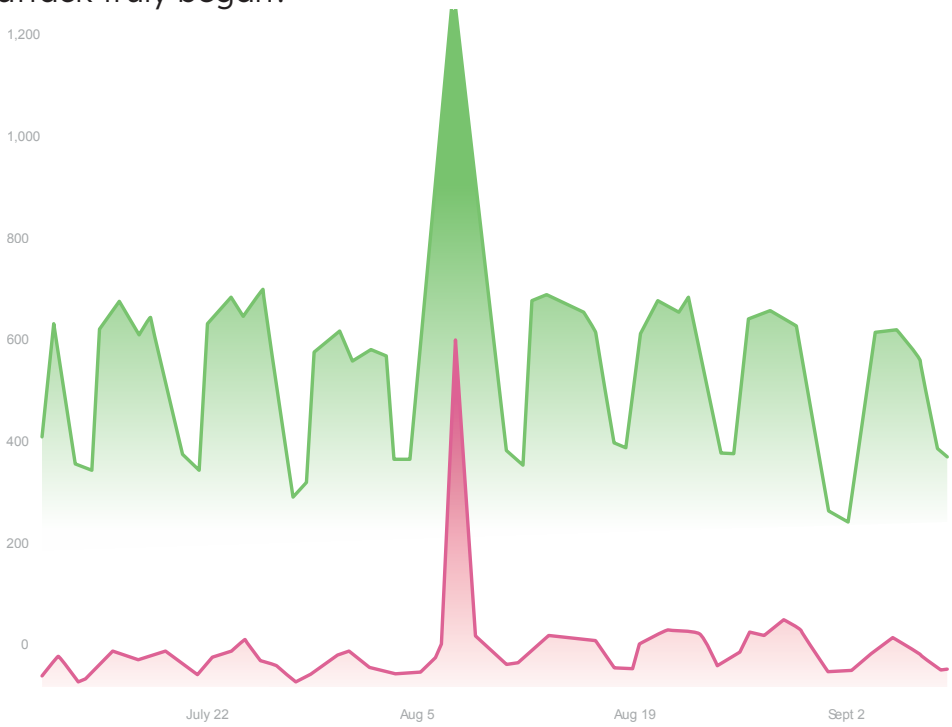


2 | While it may seem like the attack begins the week of Aug. 5, it starts the month before with a slowly escalating human force deploying a bot attack to deter suspicion.

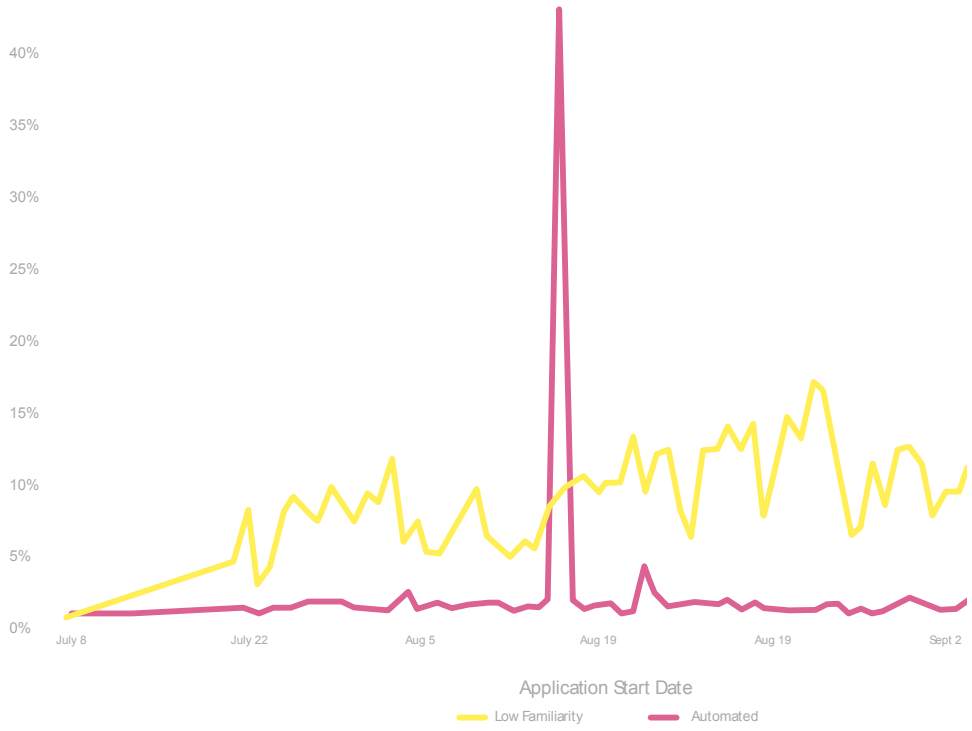




neuroID
A part of Experian

In scenarios like this, quantity gives a tactical edge — a massive distraction at one entrance will make it easier for a successive wave of fraudsters to infiltrate through another. If your manual review team is busy warding off the bot battalion, the next unit of special forces fraudsters is more likely to sneak through.

## Countermeasures

A rise in risky traffic is always worth paying attention. NeuroID is unique in our ability to separate our client's pre-submit data, making it easy to observe even minuscule increases in human fraudsters. When application volume is in the thousands, it's easy to dismiss a few hundred extra risky users, but as demonstrated, those human fraudsters can deploy thousands of dollars and cost hundreds of thousands of dollars in damage. We alert our customers to these upticks and help them analyze the trend to blocklists the device and have their team on high alert for an attack.

This lending platform was trialing NeuroID capabilities when the fraudster's feint attack began. With our behavioral visibility, their fraud prevention team was able to see their application had been targeted and comfortably infested by fraudsters smarter than they'd give them credit for.

## Actionable Intelligence

**This feint was a designed distraction:** the human fraudsters were always the real threat, and the bots were never meant to successfully complete the lender's onboarding process. But even if this bot surge (and, for that matter, the human threat) is stopped before any fraud losses occur, the costs of unnecessary, excessive data calls and manual reviews ripple across businesses. On average, every $1 of fraud losses costs businesses $3-5[6], depending on the type of fraud and industry. The impact only grows when fraudsters leverage high-volume tactics like these — it's crucial to optimize your fraud stack to stop fraudsters before they can incur downstream costs. If that means implementing behavioral analytics to force fraudsters out early and inform downstream checks for users, it's well worth the investment.

6. Current Fraud Losses for Financial Institutions, SQL Banking Systems

# Attack Style: Adaptive Tactics

In this attack, fraudsters apply their knowledge creatively to overcome new obstacles, following Sun Tzu's advice to respond swiftly to unexpected changes and continuously refine strategies based on both victories and defeats. **Here's how you stop them.**

## Story from the Frontline: Fraudsters Learning & Adapting in Real-Time

A top 20 retail credit provider observed an increase in application volume and auto declines between December 2024 and February 2025. This was only the beginning: even though NeuroID data shows winter is traditionally a popular time for fraud attacks[7], the scale and sophistication were nothing short of shocking — more than 20K bots worth $8M+ in potential fraud losses and data costs alone.

> **"Tactics are like unto water; for water in its natural course runs away from high places and hastens downwards. Water shapes its course according to the nature of the ground over which it flows."**
>
> – Sun Tzu

7. https://www.neuro-id.com/resource/report/fraudsters-almanac-emerging-trends-series/

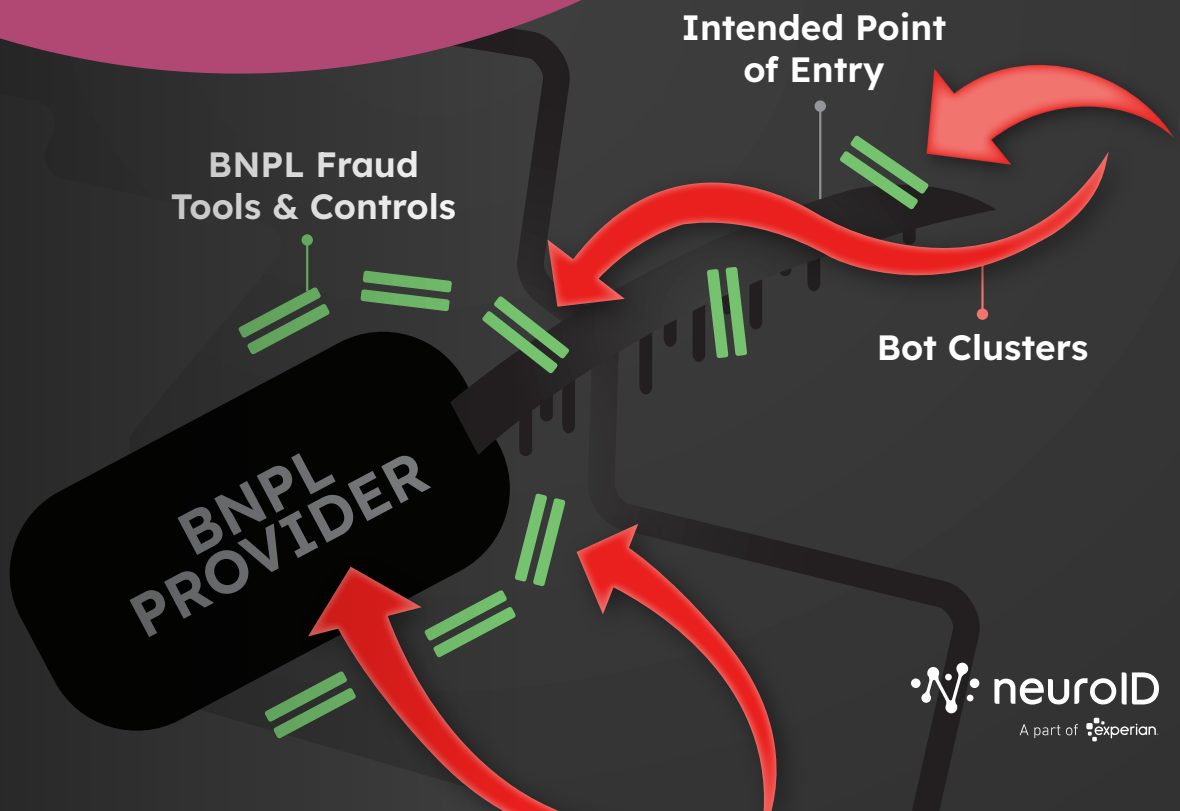




neuroID
A part of Experian

## Tactical Breakdown

NeuroID visualizes the behavioral pathways of all users through a customer's application, segmenting genuine and risky users, and further differentiating types of risky users — human fraudsters, fraud rings, and bots. For this buy-now-pay-later (BNPL) provider, NeuroID honed in on the bot traffic and identified alarming insights. More than 25K bots started the application during the 90-day period, forming four distinct clusters: two clusters (roughly 20K bots each) were deterred by the initial phone and email requests or one-time-passcode verification (don't be relieved just yet because they were deterred — there are still serious issues with these bot armies) and two clusters (roughly 3.5K bots each) that successfully completed the application (the other 1.5K abandoned the form in much smaller clusters).

Before reviewing each cluster, it's critical to understand **this bot army's approach is like nothing we've ever seen before.** Historically, we've seen clusters form sequentially, where one cluster learns the application and drops out without completing it. A new cluster returns days or weeks later and makes it a little further. Another better-trained cluster follows, and the cycle continues until one is able to successfully complete the application. **But in this case, all four clusters are present throughout the 90 days, a new pattern revealing real-time adaptive tactics NeuroID hasn't studied before. These concurrent clusters suggest GenAI is advancing fraudster tactics even faster than we thought.**

## Stages in the BNPL Loan Application

A typical BNPL loan application starts on a merchant's website. A prompt directs the user from an item to the provider's application, which includes a one-time passcode and requires the user to provide their name, address, birthday, email, and phone number. Below is what we'd expect to see from 25,000 genuine users, based on the provider's historical data.

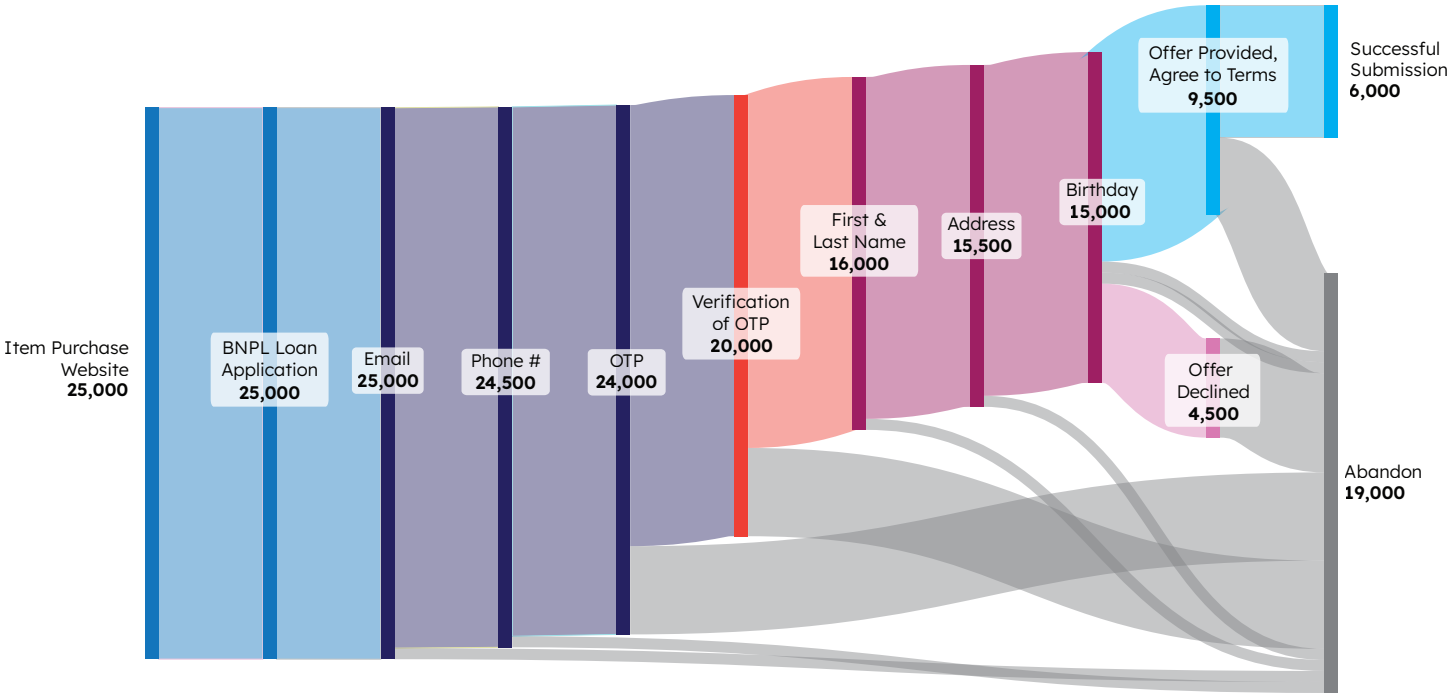




neuroID
A part of Experian

## Step Paths for Bot Clusters

Of the 25K+ bots that started the application over the 90-day period, over 3K successfully completed the application. While OTP successfully deterring 20K bots my seem like a win at face-value, there are problems:

1. OTP data calls are usually cheap, but still not free.
2. The flood of volume, as discussed in earlier tactics may have no intent of completing the application. They may purely be present to create a technical system overload or manual system overload.
3. Any fraudster spending time in your application is bad. New-aged bots are incredibly skilled in mapping your flow and returning successfully with a larger army.

Let's review both the bots who didn't get through and some of the 3K that did.

---

### Concurrent Cluster 1A
(**Fail:** dumbest – 15K)

**15K bots** enter the form as expected: through an item of purchase on a merchant's site. About 2K drop off when shown OTP (deterring as designed) and the rest drop off after failing to complete the OTP. A win? Yes and no. You successfully kicked them but at what cost? While OTP is traditionally inexpensive, it's still an unnecessary cost for 12.5K bots. Worse, most of these applicants were from a lead generator. Often, volume is valued over quality from lead generators. But as AI dramatically increases the scale of fraud, this example highlights the importance of quality assurance from your vendors.

### Concurrent Cluster 1B
(**Success:** smarter – 2.5K)

**Roughly 2.5K bots** successfully complete OTP verification and proceed to submit the application with sufficient PII data — name, address, and birthday. This group will now take up manual review hours and could potentially be offered a loan. If they're not stopped, it could cost the provider $5.5M in fraud losses.

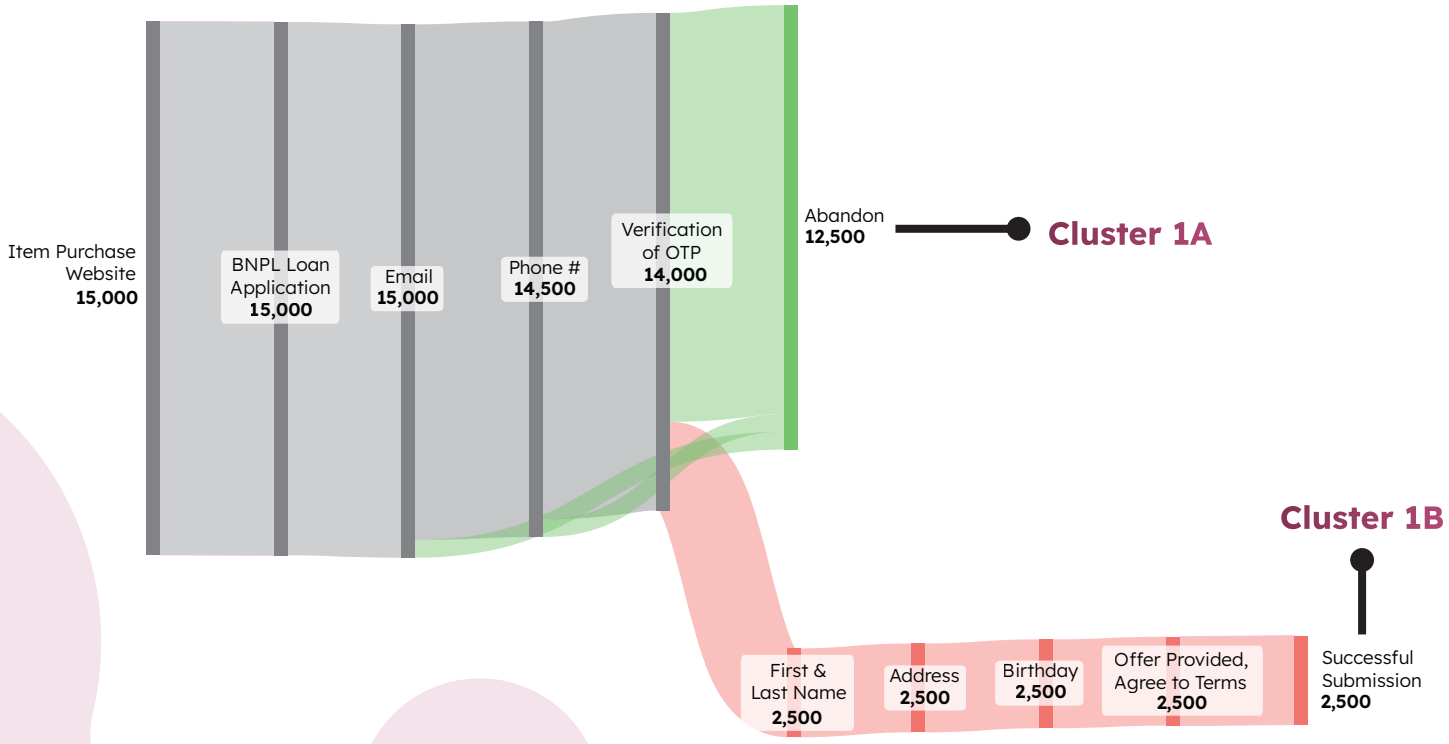




neuroID
A part of Experian

## Concurrent Cluster 2A
(**Fail:** smart – 10K)

**10K bots** enter the form through an unknown path. This is a major concern — rather than entering the loan application from an item of purchase — the foundation of a BNPL loan, these fraudsters have found a back door and are exploiting it. However, the majority of this group is successfully deterred by the OTP. Although they failed, they still demonstrated a successful exploit of a vulnerability.

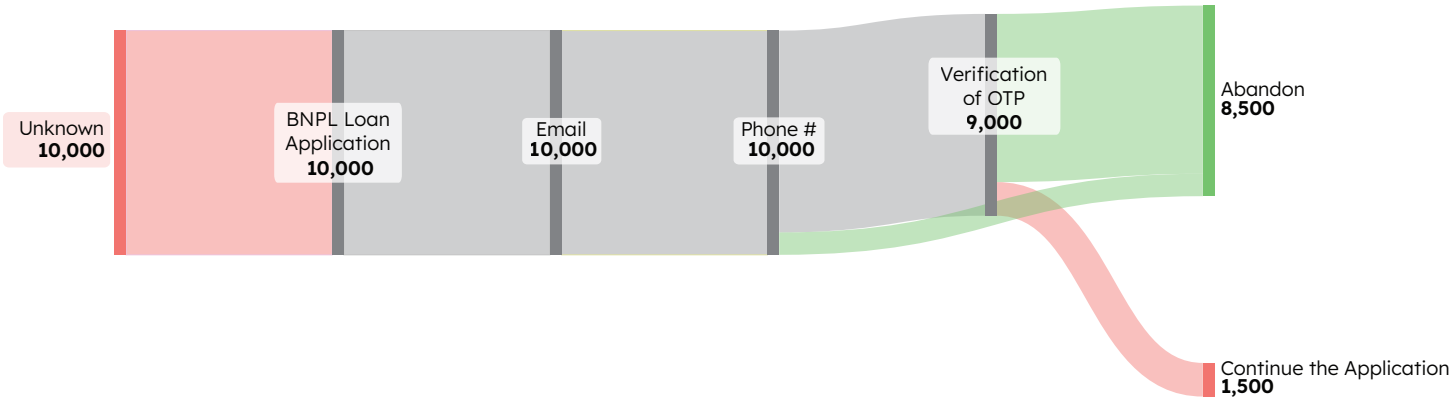


## Concurrent Cluster 2B
(**Success:** smartest – 350)

**350 bots** exploit the same vulnerability as Cluster 2A, skipping the item of purchase and directly entering the loan application. However, they also successfully bypass providing email and phone and submit the application with sufficient PII data. Although this group is small, the fraud losses could cost up to $700K for a BNPL of this size.

Skipping steps in an application is an even bigger red flag than entering from an unknown location. We see this with other customers, and while they know it can and does happen, NeuroID helps them see which users and their devices are doing it as well as exactly what fields and checks they're circumventing.

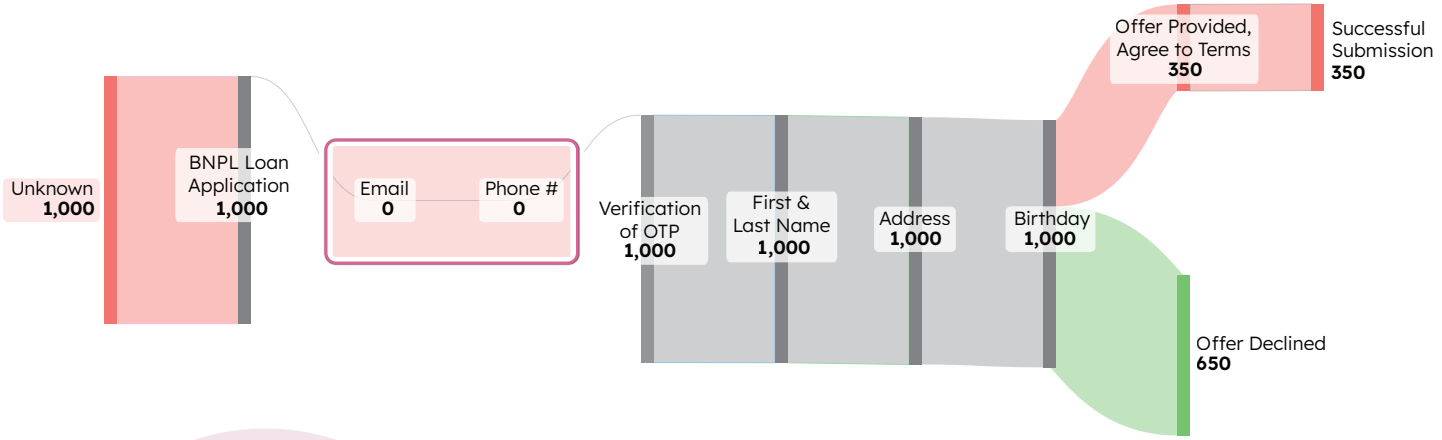




neuroID
A part of Experian

## Countermeasures

Once NeuroID revealed the infestation of bots, the credit provider quickly changed their practices to shut down bot traffic earlier in the application. They began an investigation into their lead generators and found that a significant portion of the provided leads were fraudulent. They also investigated vulnerabilities in their form that allowed bots to bypass steps and fields. This customer has a new approach to fraud. They're now reviewing application traffic regularly to identify new fraudster clusters — allowing them to adapt tactics as quickly as their attackers.

## Actionable Intelligence

These four clusters could be multiple independent fraud rings, but we're inclined to believe this scheme is more sophisticated than that. It's possible the human commanders of these bot armies are testing the provider's countermeasures, running different bot scripts depending on which applications are declined or reviewed that week. By adapting and rotating through different approaches, they gave our customer the appearance they were stopping enough bots to prevent a deeper investigation.

However, there is a more alarming possibility. Could the individual bots be learning the most effective way through the application in real-time? Each bot may be leveraging machine learning to evolve and iterate as they progress through the application. Disconnected from a larger hive-mind, could each failed bot be repeating the application until it independently finds the path of least resistance, like water running through rocks until a canyon is formed?

Fraudsters have always evolved, but this adaptive, split-second learning powered by GenAI has created a new type of warfare. Fraudsters don't need to analyze data or wade through paperwork before unleashing new tactics or tech: their attacks can improve mid-strike. This can be overwhelming, but with real-time insights from NeuroID's behavioral analytics, you can launch counterintelligence initiatives to stay ahead of fraudsters. With behavioral insights, you quickly close exploitable gaps, and send a clear message to your opponents: **attacking here is not worth your trouble.**

## Behavioral Deterrence: Stopping Fraud Before It Starts

### Allows us to ask key questions

These countermeasures are proven to help win battles against fraudsters. But to win the war, a strategic recalculation is in order. For over a decade, NeuroID has worked alongside our customers to analyze attacks and uncover the latest in fraudsters' strategies. With data on third-party fraud attacks of all kinds and nearly a billion digital events under analysis, we encourage businesses to consider:

**Q:** If you have a high number of users starting an application but a smaller number finishing, why are they dropping off?

**Q:** Is the cost of deterring fraudsters higher than it needs to be?

**Q:** Why are you attracting so much fraud, regardless of how much you can deter? Can that volume be decreased?

**Q:** Are fraudsters entering your application or circumventing key components?

**Q:** Do you have applications suspiciously incomplete?

**Q:** Do you have visibility into how fraudsters are manipulating your process to fast-track their applications?

**Q:** What interventions can you introduce at high-risk points without disrupting genuine users?

Armed with GenAI-enhanced tools and next-generation bots, fraudsters outman — and out-plan — even the most sophisticated fraud stacks. Today's fraudsters attack like stealth fighters, flying under the radar of traditional fraud detection tools and powering through makeshift barriers put up by unprepared businesses.

NeuroID is a spotlight that reveals the fraudsters lurking in the shadows. Behavioral data unveils fraudsters' battle plans and reveals a new level of opponent intel that other tools can't. With NeuroID, businesses can see and stop attacks before they even begin.

For more on how NeuroID can protect against evolving attacks, **visit NeuroID.com.**

## Footnotes:


1. Who is winning in the race for GenAI weaponization: fraudsters or fraud teams? NeuroID
2. Sun Tzu, The Art of War
3. New Industry Report: Are Fraud Bots Beating Behavioral Analytics? NeuroID
4. 5 Fraud Forecasts: What We Got Right and Wrong in 2024 So Far, NeuroID
5. New Industry Report: Are Fraud Bots Beating Behavioral Analytics? NeuroID
6. Current Fraud Costs for Financial Institutions, SQL Banking Systems
7. The Fraudster's Almanac, NeuroID