The Threat of Serial Fraud

How to Combat Financial Crime in Document Automation

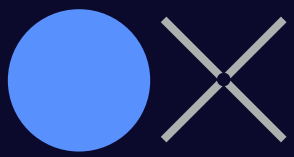

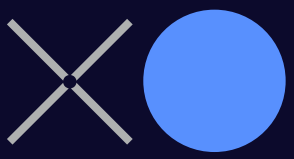

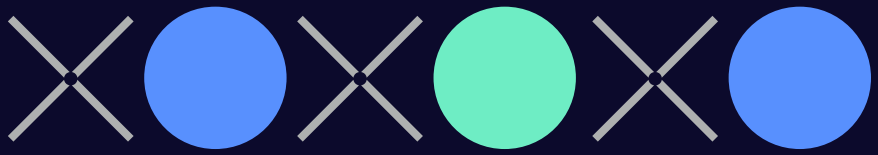

resistant ai

A new era in document fraud

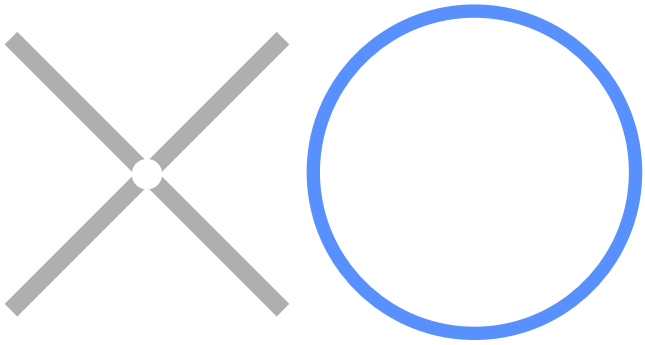

A new era in document fraud

The tritest comment in the last 25 years has to be: “Technology is changing everything.” But that’s just the symptom of an acceleration that is hitting an exponential curve. Case in point with financial services. Open an account from your living room in minutes, send money in just a few taps of your screen, or apply for a loan that’s personalized to your true credit profile. Services that just a few years ago were still considered science fiction are now everyday realities—and all of them rely on some form of automation to achieve scale.

Technology is neutral, though, and an equally dramatic transformation has been happening in fraud and financial crime.

No longer do criminals need to risk committing crime in person: they can do so behind the safety of a fake identity, operated remotely from a jurisdiction on the other side of the globe with no extradition treaty. And what used to be small opportunities can now be reaped at scale thanks to automation: Opening a fake account to take advantage of a bank’s free $10 promotion might not seem worth the effort, but opening 100,000 accounts in three days with scripted automation fundamentally changes the equation.

The place both sides intersect is, in some ways, the most archaic part of customer verification: documentation. What was once physical is ported over into the digital realm, creating opportunities to onboard customers with unparalleled customer experiencs—and the risk of onboarding fraud at a scale never seen before.

And scale is what defines serial fraud.

If the initial phase of the digitization of financial documentation brought about document fraud that could no longer be detected by the naked eye, serial fraud represents their industrial (re)production and distribution via automation practices.

> **Serial fraud is the industrial (re)production and distribution of fraudulent documents via automation practices.**

In the context of online financial services, serial fraud is the systematic reuse of a fraudulent ID, proof of address, proof of income, or other KYC/KYB document, with changes in details done at high volumes—essentially turning that document into an editable template for more fraud. In their most severe form, these templates are created and distributed via automations that allow single individuals or organizations to scale their activities.

Serial fraud takes two different forms:

Decentralized serial fraud

Retail fraud-as-a-service models distribute editable fraud templates via search, social media, and marketplaces at low cost to a large segment of end users who lack the technical capability to create fraudulent documents themselves. Typically, they feed large volumes of first-party fraud targeting online lenders and tenant screening efforts.

Usually, these services operate as online PDF editors whose practices can be modeled relatively easily given enough samples. There is little incentive for the producers to update these templates, and they can stay in circulation a long time—although more are branching into different document formats.

Concentrated serial fraud

Here, enterprise fraud-as-a-service models with startup-like operations leverage automation and iterative experimentation to test and bypass automated controls, providing organized crime rings with unprecedented scale or reach. Typically, the goal is the mass creation of accounts, which become commodities that can be leveraged or resold for various financial crimes, from coordinated bust-out frauds to money muling.

Unsurprisingly, this more sophisticated form of serial fraud is commonly found in payment providers, banks, and other financial services or marketplaces that offer accounts. Given the stakes, the documents used are often brand new leaks or productions that require comparative techniques that go beyond simple document modeling.

The threat of serial fraud

Serial fraud creates two very clear challenges for financial institutions:

Document automation risks become exponential

Automating the intake of documents is one of the highest impact activities a financial institution can carry out to increase efficiency and reduce costs. Intelligent document processing (IDP) services extract the contents from documents using parsing or optical character recognition (OCR) faster and with more precision than humans, and these services analyze the output for name- or address-matching, credit scoring, and more. Many are now deploying large language models (LLMs) to contextualize the content in a way that seemed impossible not long ago.

Unfortunately, these systems are generally built with trust in mind and are incapable of detecting forged documents. This means that deploying them without a form of document fraud detection opens the door for the mass-produced templates of serial fraud to ride the automation workflows and overwhelm any follow-up review processes.

It's an automation arms race

With the essentially endless volume of fraudulent documents, getting around traditional verification methods is now a numbers game for criminals. Individual fraud-fighters might be able to catch an instance of document fraud, but it's nearly impossible for them to know that they've discovered part of a cluster of near-identical documents. Moreover, they're unable to share their findings so that others might catch the rest.

Only systems that can comparatively assess all incoming documents and associated behaviors at once can tackle this challenge. Humans still need to be in the loop, but one step removed, reviewing the total cluster of related documents.

The goal of this report is to shed light on the nature and scale of serial fraud as well as existing solutions to combat it. As serial fraud remains a new, evolving threat, it's crucial for any organization accepting digital documents for customer verification or underwriting processes to be equipped with the knowledge necessary to safeguard their systems.

What serial fraud looks like

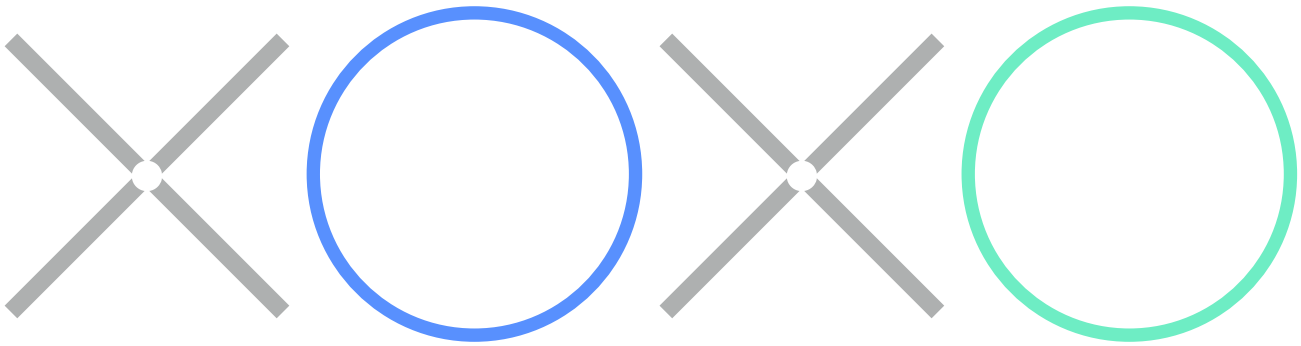

What serial fraud looks like

The golden rule of storytelling and education is show don't tell. In that spirit, we've prepared multiple examples of serial fraud, which we have carefully recreated from real detections. All personally identifiable details, including stamps, faces, and signatures, have been either replaced, blurred, or altered to protect the victims—and the guilty.

As you'll see, serial fraud occurs around the world and exists in all major formats of documents accepted by financial institutions.

GIF animation Click the image when you see this icon for an animated gif of the fraud.

Serial fraud comes in images

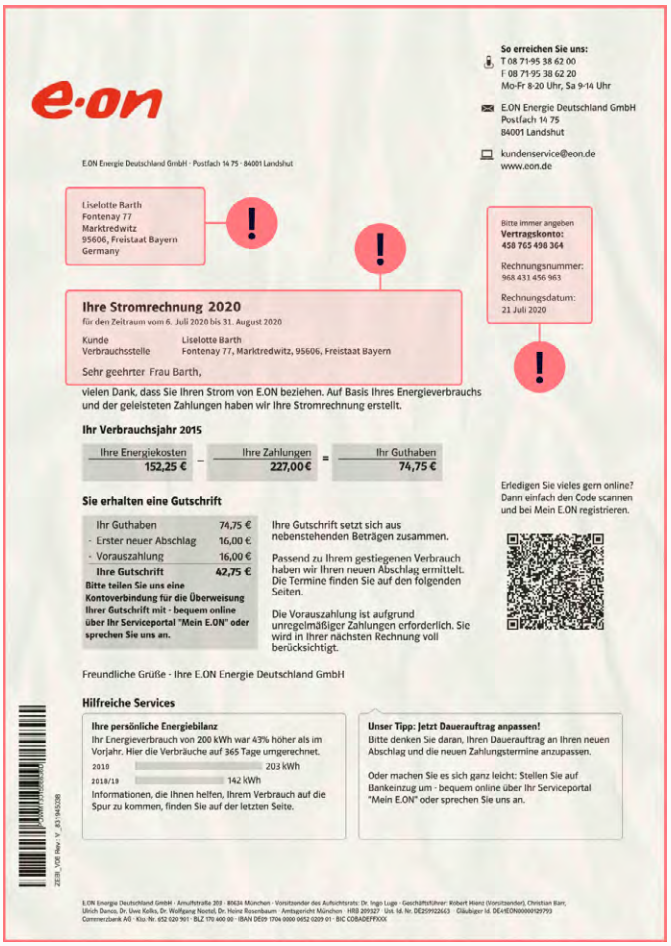

A utility bill from German energy company EON used in Proof of Address. Notice the creases and the QR code vs the details of the document changing.

resistant.ai

Serial fraud hides in Screenshots

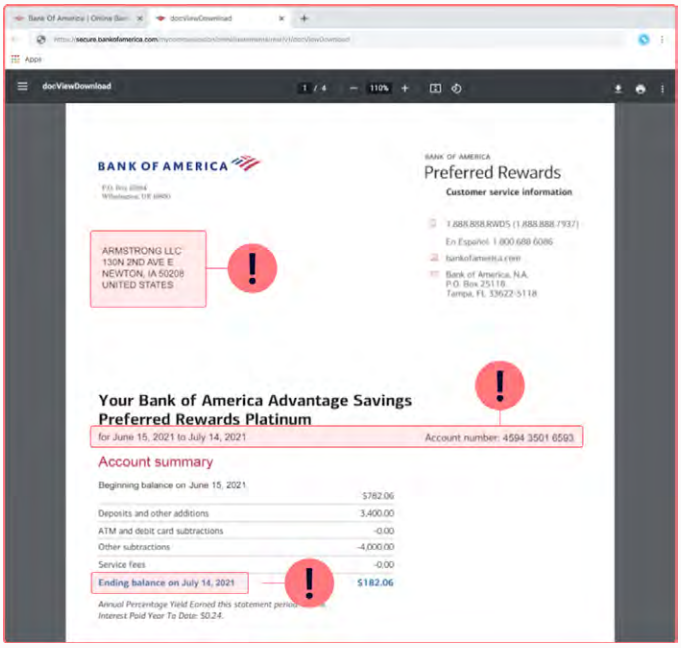

This Bank of America account statement could be easily manipulated in the browser's code, and then screenshotted over and over. Notice the URL and transaction amount stay consistent.

Serial fraud hides in Scans

This Moroccan tax certification attestation used in a merchant onboarding process was scanned and modified multiple times. Notice stamps and signatures stay consistent.

Serial fraud crosses formats

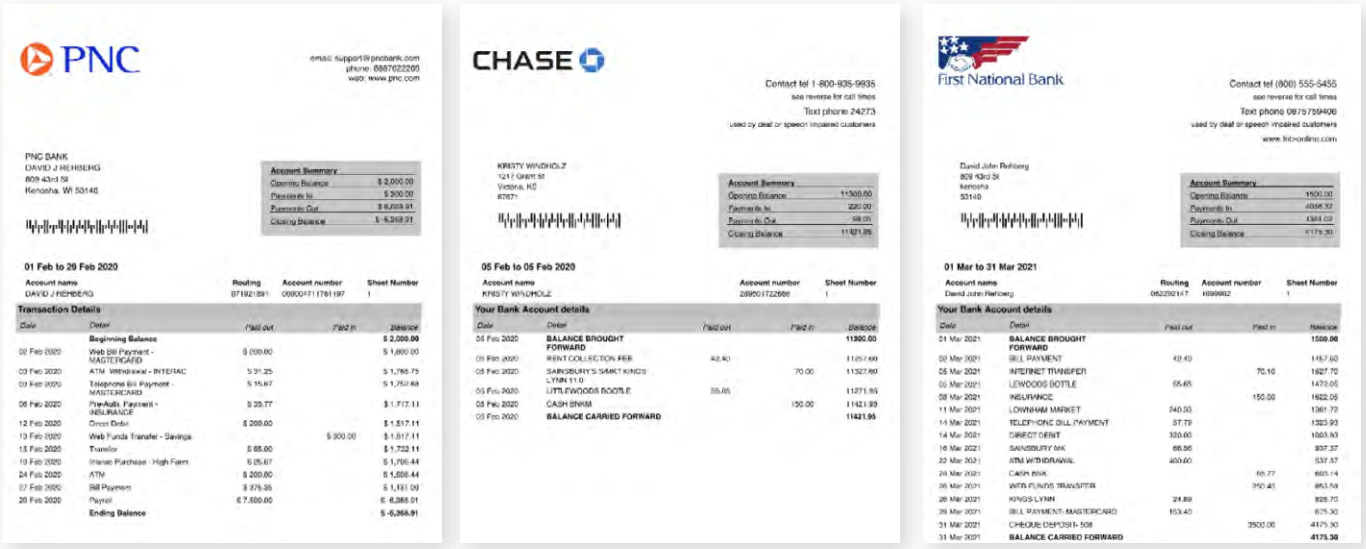

This particular bank statement template was found in PDF format in one customer's data set.

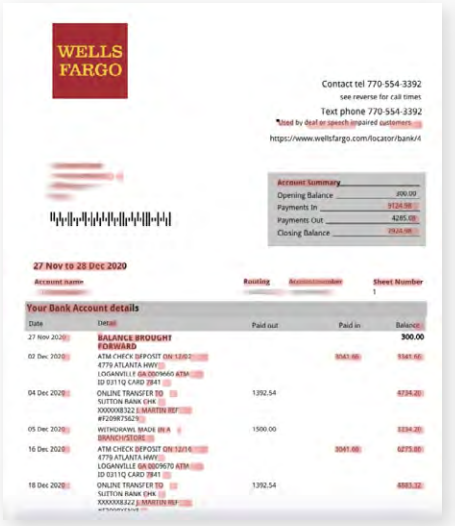

Found in a second

This capacity for criminals to leverage the same document in a wide variety of contexts not only demonstrates the almost viral scalability of any one fraudulent asset, but also the challenge in detecting them—picking up such radical transformation takes more than identifying templates.

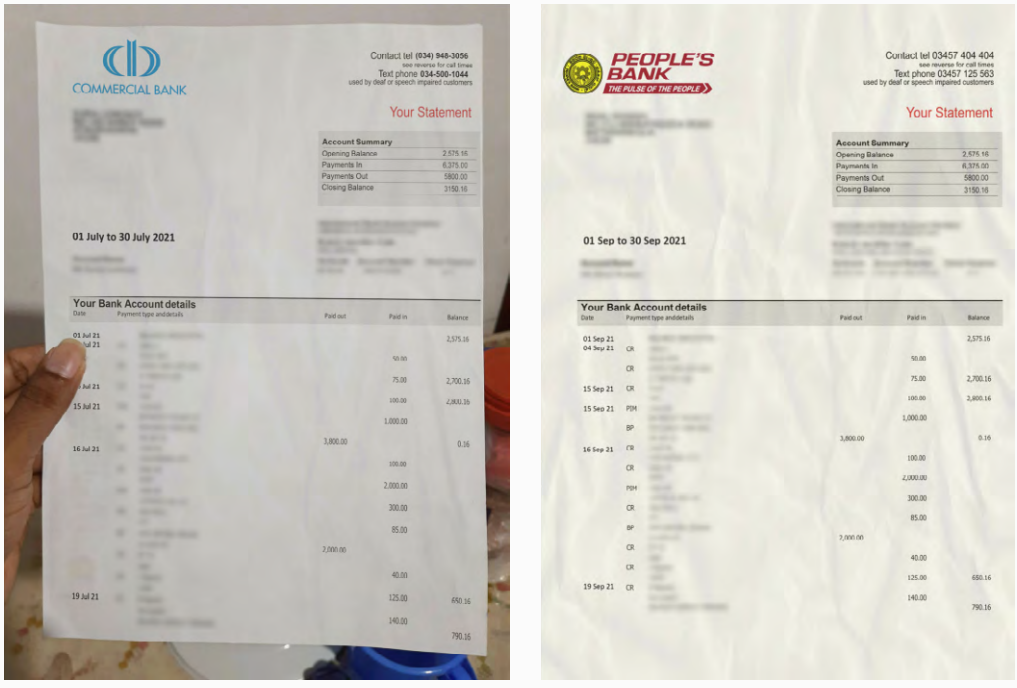

And also digitally overlayed onto crumpled paper backgrounds and placed into different contextual scenes and backgrounds.

Serial fraud uses Generative AI

Serial fraud also occurs with other types of documents used in customer verification flows. In particular, as AI gets better at generating faces, more and more serial fraud is occurring with ID documents and selfie images:

Belgian passport with generated faces

Serial fraud also includes stolen, non-fraudulent documents

The case below shows a series of stolen Mexican IDs being submitted by the same criminal in an almost industrial fashion. Notice the criminal's ring on the hand holding the cards.

Serial fraud also includes selfies stolen from social media

A common tactics is to steal a selfie from social media, and then add hands and a blank piece of paper which can be repeatedly updated.

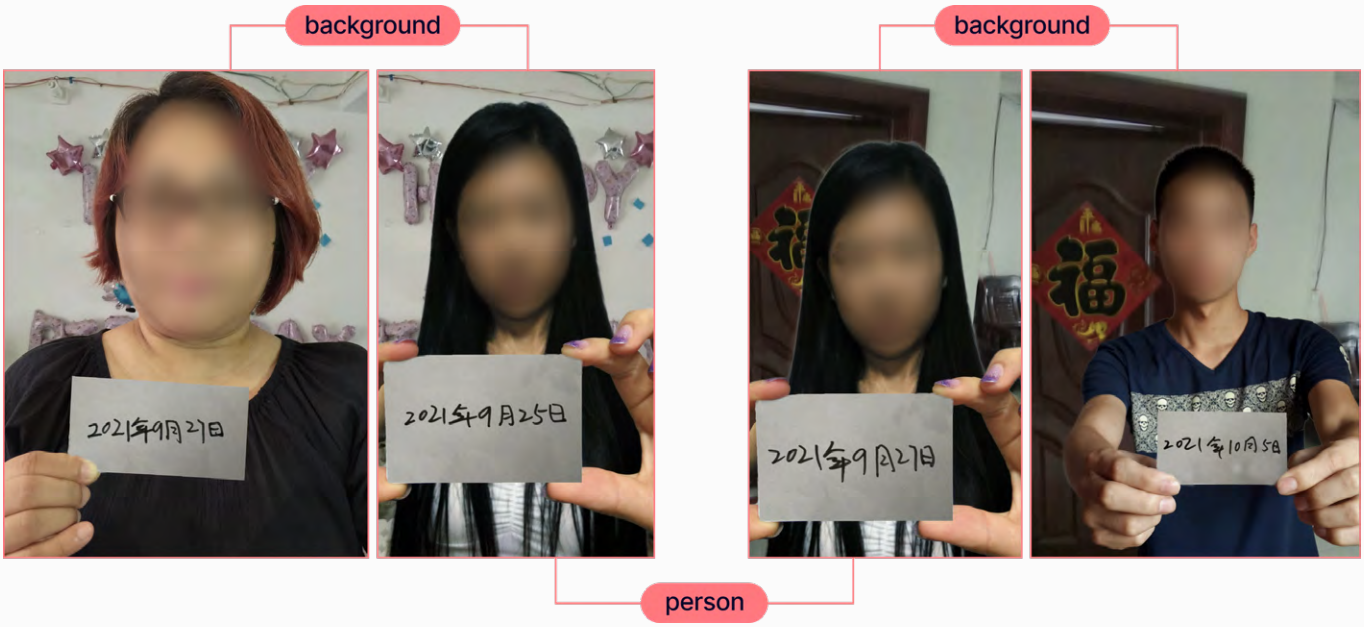

And different elements of the scene composition can be mixed and matched to make sure there is enough variation to fool automated detection systems.

The scale of Serial Fraud

Resistant AI has analyzed over 50 Million documents, and the following data represents the scale of serial fraud detected in our dataset between January and November 2023. When interpreting the data below, it's important to note how the following two concepts interrelate:

Fraud template: a document used as a template by fraudsters to create variations for submission.

A cluster: a grouping of fraudulent documents tied together either by the use of a common template or by other forms of associations, such as using the background or being submitted from the same device. A cluster can include multiple templates.

~2%
Share of serial fraud documents in all KYC, KYB, underwriting and screening use cases

While serial fraud does occur in full digital PDFs, it is much more prevalent in scanned PDFs, digital print PDFs (when an image is saved in PDF format), and regular image formats such as jpegs, PNGs, etc.

~5%
Share of serial fraud in image-based documents

>3,000
Different serial fraud clusters detected

>70,000
Most serial fraud documents detected in a single customer

13,210
The most documents in a single cluster

2,556
The most documents created from a single template in a month

Unique serial fraud clusters detected per month

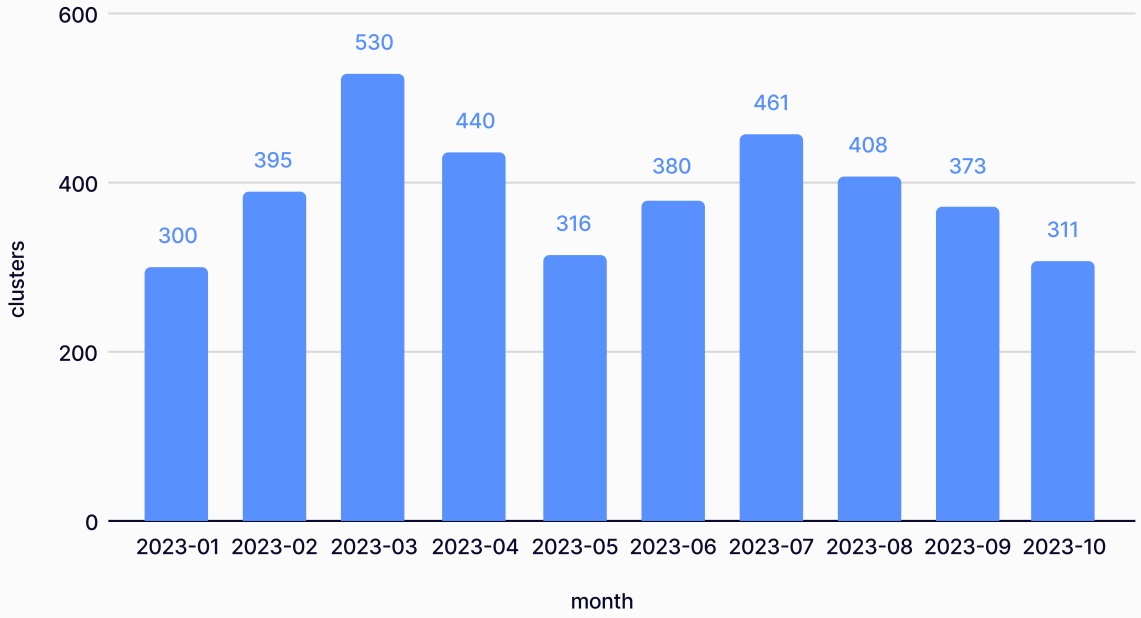

Resistant AI constantly detects >300 unique serial fraud clusters each month—usually 3-10x the findings of competing solutions.

Distribution of clusters by number of documents

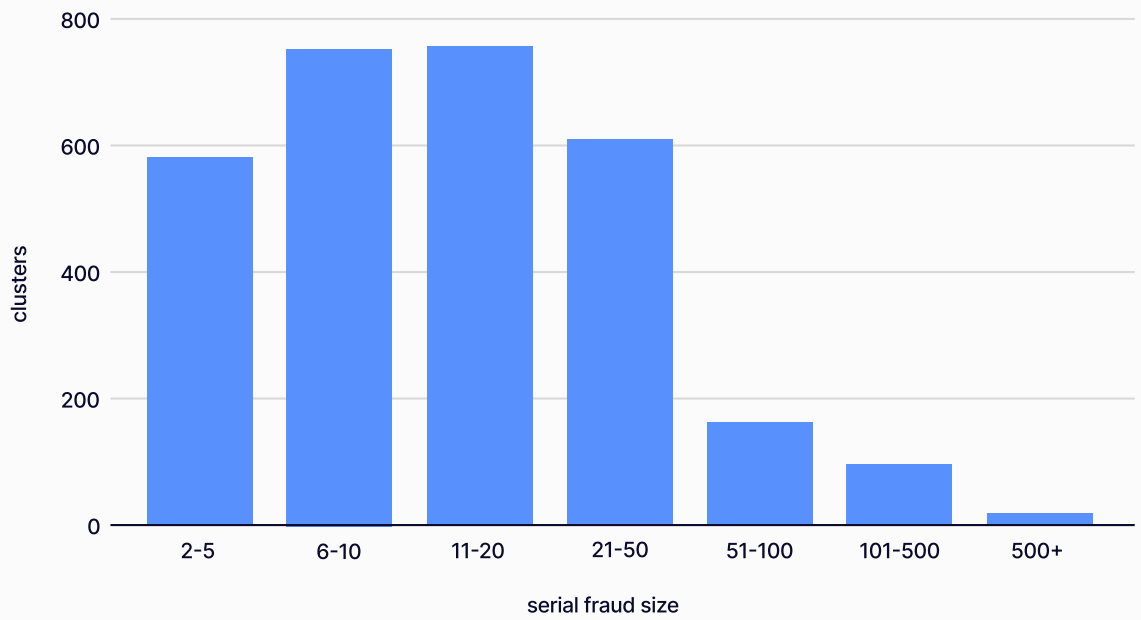

70% of serial fraud clusters are used to create less than <20 documents.
The very last 0.16% cover clusters that were used to submit >500 documents. These usually represent the outputs of online template farms, but the occasional high volume custom job is also included.

Distribution of clusters by fraud campaign length

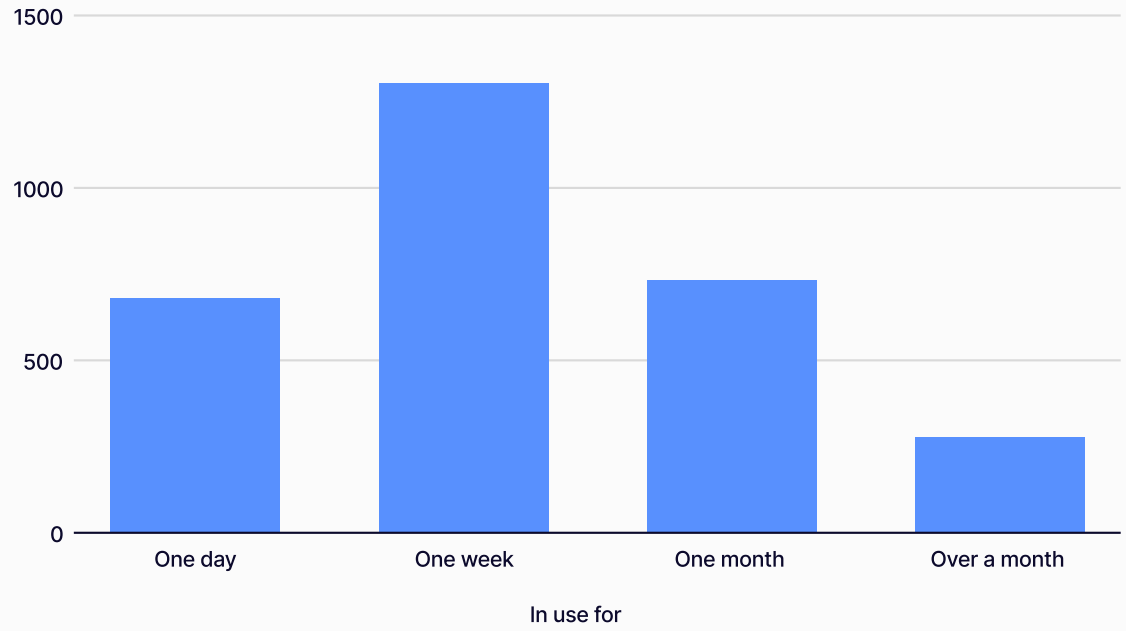

Just 30% of serial fraud events last longer than week, and they tend to be composed of online template farm outputs. Hidden in the 70% that occur over the course of a week are high-volume attacks from organized crime rings.

Profile of a high-frequency targeted attack with a Canadian passport

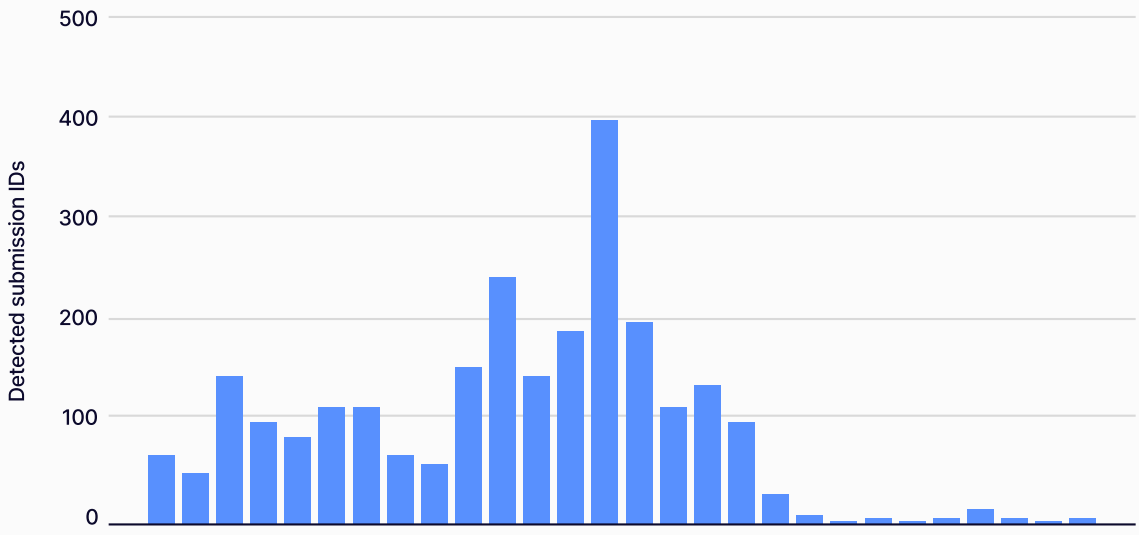

One extreme example is that of this Canadian passport which was modified and submitted over >2500 times to different customers over the course of just 20 days.

Who produces serial fraud?

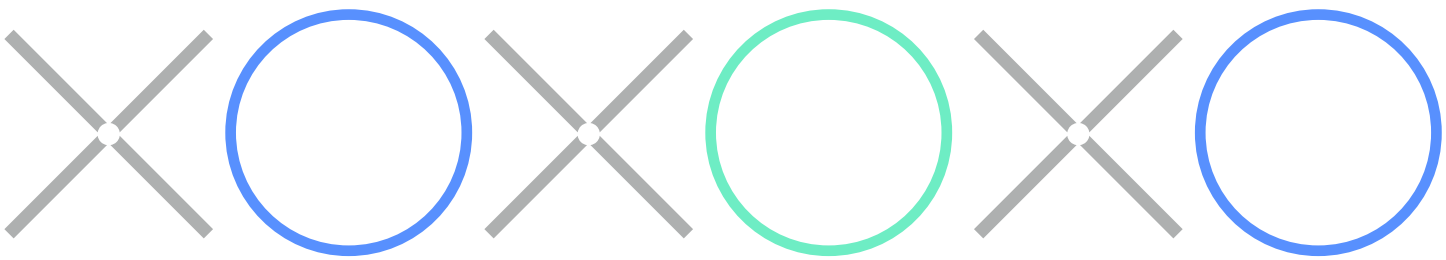

Who produces serial fraud?

Serial fraud perpetrators fall into two broad buckets, which reflect different tendencies in digital technologies: online template farms that feed the "long-tail" of potential fraudsters, and concentrated, high-impact campaigns perpetrated by well-resourced, organized criminals.

Online template farms: serving the long tail

Committing document fraud has become more accessible than ever. Gone are the days of risking your own documents, finding unique ones on the dark web, or developing the skills and know-how to produce passable forgeries. Today, high-quality, ready-to-edit templates are just a simple search away thanks to online template farms that enable customers to commit income fraud. These templates allow customers to get better interest rates or access to financial products from which their risk profiles would otherwise exclude them.

And it's big business: There are now hundreds of template farms providing templates for as low as $5 and as high as $100, and in some cases, bundled or offered as subscriptions. While some boast as many as 8,000 document templates available (fig.1), the ecosystem has grown so large that some are specializing in different document types, such as utility bills, paystubs, credit reports, tax forms, and so on.

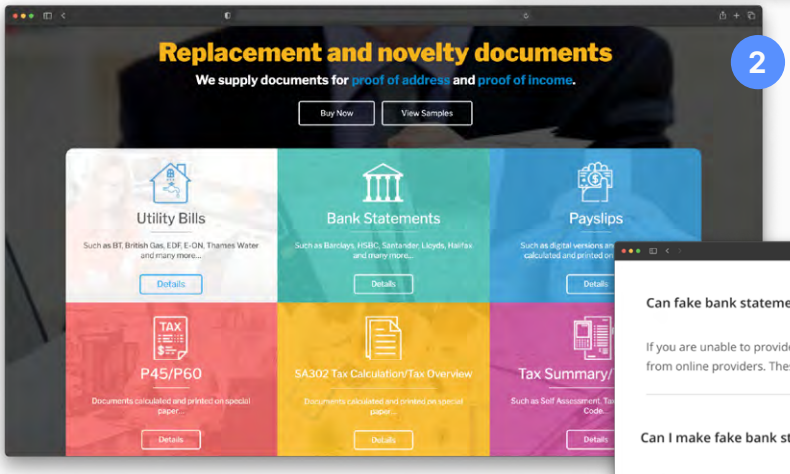

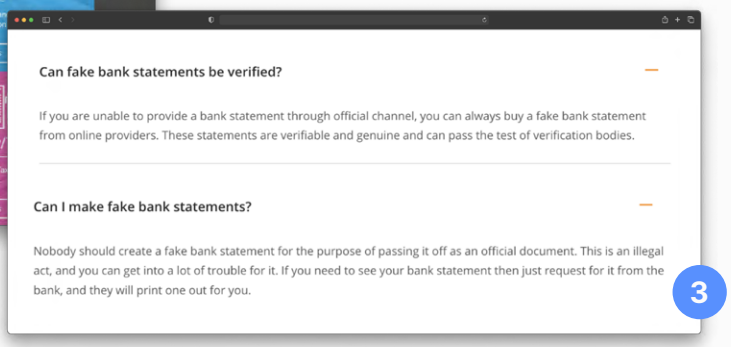

The perpetrators are fairly brazen about their activities. Many clearly lay out financial use cases such as "Proof of address" or "Proof of income" (fig.2) while also claiming that their offerings are for "novelty or personal use only." Meanwhile, FAQs reassure users their documents can pass verification in one question, while admonishing anyone that would dare consider creating fake documents and passing them off as real in the next (fig.3).

This erratic messaging is not only an attempt at (weak) legal cover, but it's also a means to hit the right search engine optimized (SEO) keywords to generate traffic: the largest sites boast well over a million monthly site visits, but even modest ones can garner 15,000 visitors per month.

Site	Monthly visits
Templatelab.com	1.7M
Stubcreator.com	55K
PSDlife	40K
Replaceyourdoc.com	18K
Onlinenoveltydoc.com	8K
banknovelties.com	3.5K

A small sample of popular template farms and mothly traffic

What's more, things are not always so clear cut when it comes to legality. A whole class of legitimate businesses unknowingly assist criminals: for example, some payroll or accounting software companies happen to offer document templates for small businesses as a marketing tactic to generate traffic for their site, which are just as easily repurposed by criminals to create legitimate looking documents.
As is typical with other forms of fraud, messaging apps, social media platforms (fig4), and online marketplaces have similarly been used to advertise, sell, and distribute fraudulent products.

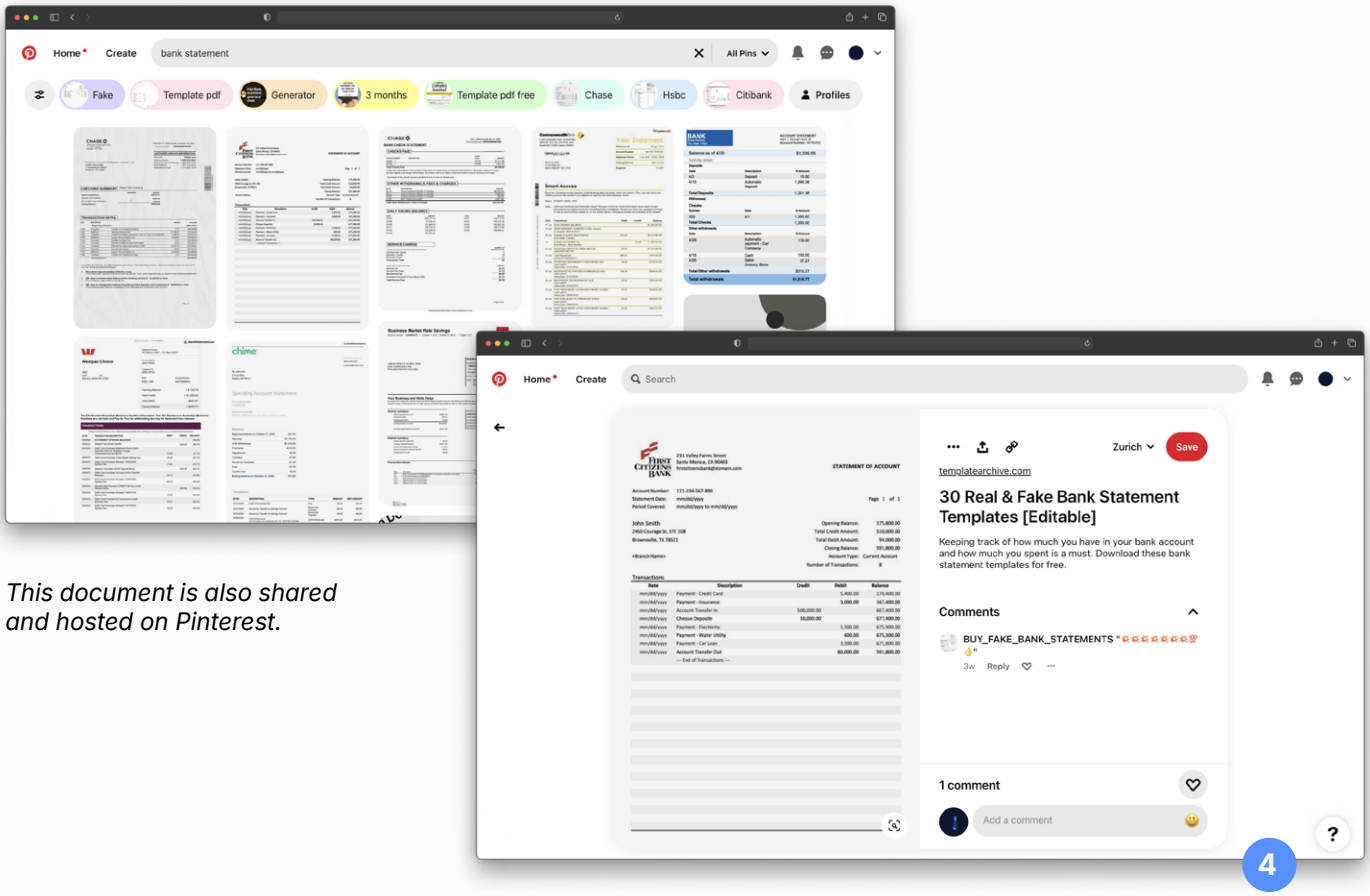

This document is also shared and hosted on Pinterest.

Fraud scaleups: high-impact campaign runners

If online template farms are the retail version of fraud-as-a-service, then organized crime are the enterprise-level operations that leverage the full scope of techniques and suppliers the underground economy can provide. This is the province of third-party fraud that has much more ambitious goals than simply pulling off income fraud for better interest rates on loans. Instead, these attempts can the creation of synthetic money muling accounts can be used for a variety of crimes:

- Large-scale promo abuse on financial services marketing offers
- Money laundering
- Coordinated bust-out fraud (either credit card or loans-based)
- Receiving illicit funds from authorized push payment scams
- Any other scheme where accounts are the enabling vehicles

There are several advantages to creating these kinds of synthetic accounts:

1. They can be created extremely cheaply and iterated upon until the weaknesses in the risk control mechanisms of the targeted institution are well understood.
2. The longer they exist, the more uses they can be put to.
3. If they are discovered, there are no real-life go-betweens the authorities can identify to work up the chain.

How these fraudsters source their documents

To substantiate these synthetic identities, these organized crime rings will rely on newly leaked, stolen, or custom-made documents in order to get around many of the template-based detections employed by the low end of document fraud detection solutions. These come from a variety of sources, such as:

- **Dark web leaks:** A variety of documents can be acquired relatively cheaply online, for as low as $9 to well over $3,000 depending on the quality and rarity.[1]
- **Individual producers:** Professional forgers still exist that haven't yet opened a store front on the web, and many have signature traits. For example, we detect over 300 documents each month in a variety of customer data sets that bear the signature traits of an individual forger based out of Russia.
- **Receptionist scams:** More and more hotels and building lobbies are becoming the source of photographed IDs taken at reception. These are then paired with forged documents to get past financial services.
- **Recruitment scams:** Whether it's remote or in person, job applicants are occasionally fooled into submitting their ID documents and/or financial documents as part of their application process.

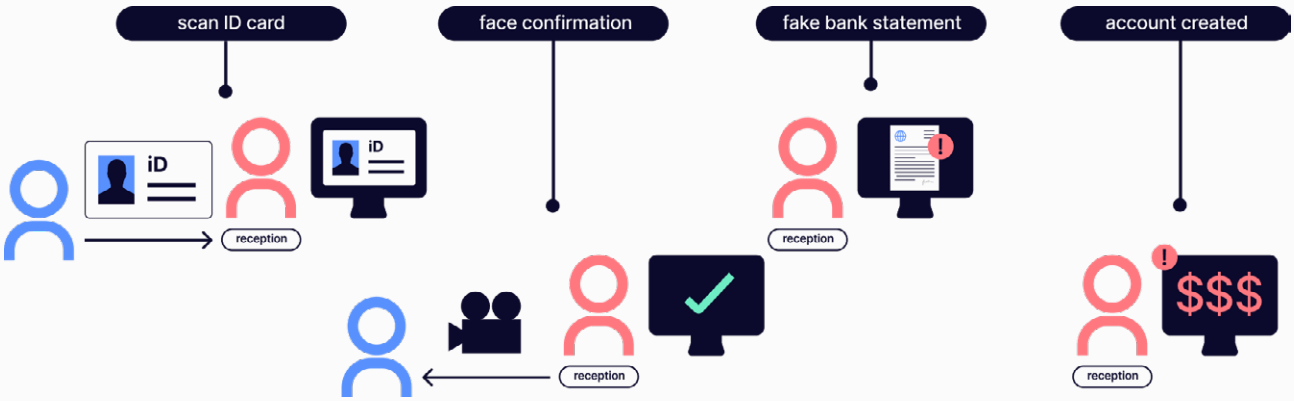

1 "Dark web market case study - NordVPN." https://nordvpn.com/research-lab/dark-web-case-study/.

How to stop serial fraud

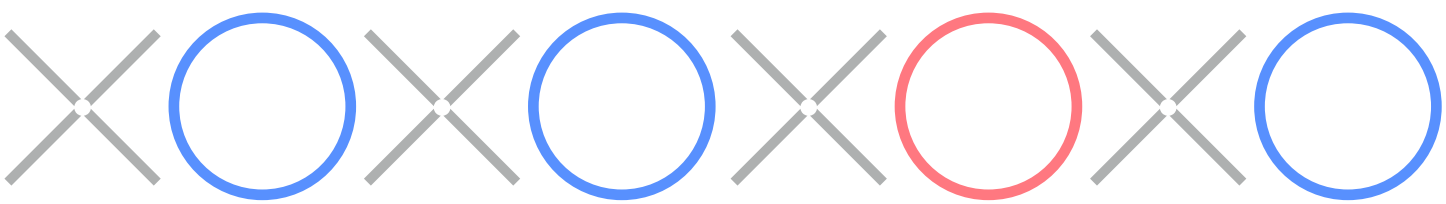

How to stop serial fraud

Some would argue that digital onboarding—and the increased volumes of fraudulent identities and scalable attacks that come with it—would require financial services to increase the friction and selectivity of their risk processes by several orders of magnitude just to keep the risks on the same level. This is not so: instead of increasing friction, existing data can be used to address scalable risks far more effectively.

This is exactly the game we play on behalf of our customers. The job of Resistant AI is to reliably defy scalable attacks using machine learning techniques that outsmart the criminal's technology.

To do so, we don't rely on silver bullets; in the real world, no technology is bulletproof. Inspired by cyber-security doctrine, we rely on layers of interlocking defenses, where multiple independent layers progressively reduce the level of threat and cover each other's blind spots, increasing the complexity (and the costs) for criminals to outwit each layer.

The Resistant AI approach: Defense in depth

Document fraud detection in individual documents

Intake customer documentation as you would normally, and our Document Forensics automatically analyzes them for forgeries. We combine over 500 different detectors that largely fall into two categories:

- Generic fraud detectors that look for signs of modification on any document—whether we've seen it before or not—such as structural and metadata anomalies, editor traces, or changes in fonts mid-word.
- Template-based detectors that know exactly how specific issuers (legitimate or fraudulent) create their documents, and whether the document matches or deviates from their practices.

This layer is already more than capable of detecting digital PDF serial fraud templates, as it only requires 20 examples to work out a new template, and delivers **uplifts in fraud detections of up to 32% while reducing manual reviews by up to 92%.**

Serial fraud detections: comparative analysis of all submitted documents

In the next layer, all available properties from submitted documents—metadata, structure, format, visual similarities, and even our own fraud analyses—are compared against each other to detect repeating fraud patterns, document reuse, stolen documents, and other attempts at high-scale fraud. Things as obvious to the human eye as repeating background scenery to as subtle as similar repeating color profiles or resolution levels are analyzed with anomaly detectors to pinpoint and flag statistical anomalies for review.
This layer is especially effective against image-based serial fraud attempts, including pre-digital modifications, whereby fraudulent documents are printed and then photographed or scanned, and can produce **uplifts of up to 98% on top of the previous layer.**

Unknown documents

Create connected components based on a single–document features (image hashes, image properties,...)

Detect fradulent components with contextual features (self–similarity, component differences)

Suspicious submissions detections: behavioral and identity analyses

Any additional behavioral data—from basic server logs to device intelligence, along with the personal information submitted as part of the application processes (names, emails, phone numbers, forms, etc)—is used to detect repeated or automated requests to create accounts. These are clear indicators of a bad actor trying to penetrate your systems.

These checks are known to **increase detections by yet another 59% on top of the combined previous layers.** More importantly, they also work to tie together what seem like separate fraud attempts into clear patterns of coordinated attack.

Ongoing monitoring detections: tying documents to transactional behaviors for perpetual KYC

In real time, our Transaction Forensics layers on top of your existing transaction monitoring system to help you to identify behaviors consistent with fraud and money laundering. Specially developed detectors use statistical anomalies to flag suspicious money flow patterns and relationships between accounts—relationships that point to fraudulent accounts, account takeovers, and other bad actors all working in favor of fraudsters.

Combining insights from customer documentation, identities, behaviors, and transactions—both for new and existing customers—prevents the creation and scaling of synthetic money mule accounts, and returns the initiative to the financial institutions. This layer has the most modest improvement on detections (14%) but also **reduces the total number of clusters by 40%**—meaning it provides you the threat intelligence to know which of the seemingly independent criminal events are actually part of coordinated fraud campaigns.

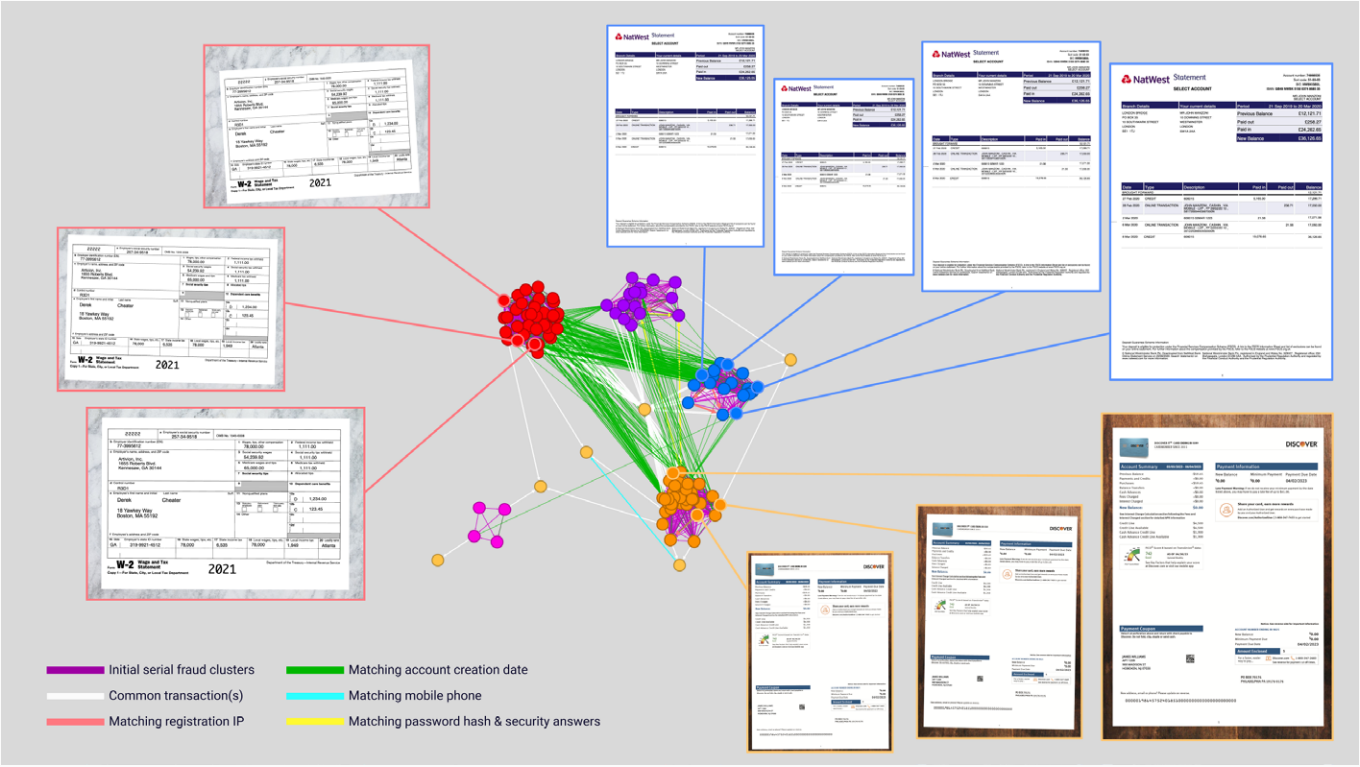

Detection layer	Cumulative uplift in fraudulent account detections
Single document fraud detection	+32%
Serial fraud document detection	+98%
Fraud submission detections	+59%
Transactional data detections	+14%
Final uplift of all layers	260%

Steps that your organization can take now

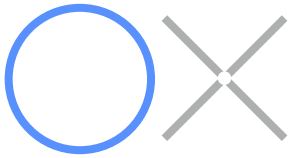

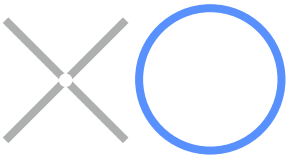

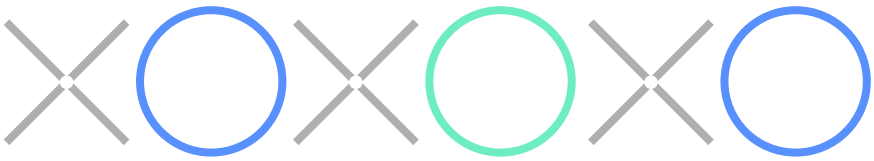

Steps that your organization can take now

Audit the documentation of your known fraud cases

While stopping serial fraud in real time is impossible for humans to do, a retrospective comparison of documents in known fraudulent accounts or applications can easily find obvious template reuse in the case of photos, or at least determine a repeating pattern of document types in PDF files. Whether you find clear evidence of serial fraud, or simply enough to warrant suspicion, it's time to move on to the next step.

Start layering AI into your document intake

The number one mistake we see financial institutions make regarding document fraud is waiting to build a full end-to-end automated document processing workflow before considering adding a document fraud detection layer.

While solutions like Resistant AI do integrate into those workflows by API, they can be leveraged manually as well to start reducing exposure to fraud campaigns and save customer approval representatives, underwriters, front line and second line investigators significant amounts of time authenticating customer documentation.

It also gives you the opportunity to build out the risk decision workflows that you can later automate.

Look for intelligent document processing solutions that take fraud seriously

Financial services face fundamentally different risks than other sectors that leverage document automation. When picking an intelligent document processor, make sure they understand those risks, and have some form of fraud prevention that is either built-in or can be integrated with.

Resistant AI already partners with many of the leading players in the market, such as Instabase, Kofax, and Google Document AI, to name a few, and can sit in front of many others in your custom workflows.

Resistant AI can help you with each of these steps.
To get in touch, please reach out to us at resistant.ai

Joe Lemonnier

Head of Product Marketing, Resistant AI

Joe has been digging into the intersections of tech and criminality for over 10 years, and brings a wealth of experience in customer research, design, product strategy and development, and content and thought-leadership to Resistant AI. His skills were honed in cybersecurity, online privacy, and productivity services that met the needs of hundreds of millions of users, and he's helped to drive business growth through service expansions and revamps, new-product discoveries and launches, and portfolio streamlinings. Engaging with customers to define seamless protective solutions are the main things that get him up in the morning.

www.linkedin.com/in/jklemonnier/

Joe.Lemonnier+bio@resistant.ai

resistant ai

sales@resistant.ai

www.resistant.ai

Credo. GV Index Ventures Seedcamp NOTION