kyndryl

Synthetic identity fraud

Understanding and combatting the fastest-growing financial crime

By Robert Wallos

Contents

PAGE

02 Introduction

03 Understanding the cybercriminal and their ecosystem

04 Why synthetic identity fraud succeeds

05 How synthetic identity fraud works

06 Why traditional detection fails

07 Steps to fight back

09 Kyndryl's perspective

Introduction

Across our increasingly digitalized global economy, fraudsters are finding new and creative ways to exploit financial systems by creating fake identities with a combination of real and fabricated personal data. Known broadly as synthetic identity fraud (SIF), this variety of theft is rapidly becoming one of the most prevalent forms of financial crimes globally.

Unlike traditional identity theft, criminals creating these synthetic identities (SIDs) are often able to pass verification checks, establish credit histories, and operate undetected for long periods of time before going on to commit fraud. And in a globalized market now-dominated by digital banking services—particularly in a post-pandemic environment—the attack surface for such fraud has expanded tremendously. That makes it harder for institutions to detect fraudulent activity. In fact, traditional fraud detection systems are often ineffective in identifying highly complex SIF operations.

SIF is the fastest-growing type of identity theft, accounting for 80% of all new account fraud.[1] It is projected to generate at least $23 billion in losses by 2030.[2] Despite its prevalence, many people are unaware of its existence.

While the industry and authorities have identified the growing problem, legacy systems and approaches cannot reduce or even slow the growth of SIF. An increase in online commerce, along with an increase in data breaches, are contributing to the growth of this criminal activity. As a result, businesses and regulators around the globe are investing in advanced identity verification and machine learning fraud detection to combat this evolving risk.

Synthetic identity fraud by the numbers

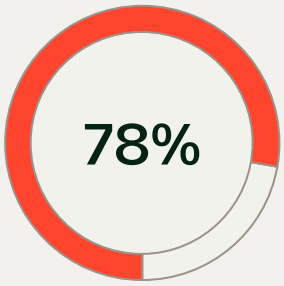

The percentage increase in data breaches from 2022–2023, contributing to the rise of SIF.[5]

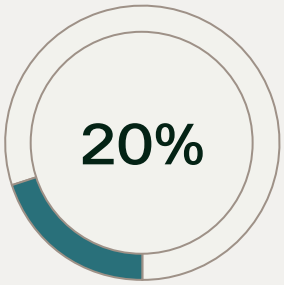

The estimated rate of all credit card losses attributed to SIF.[3]

$15,000

The average per-incident loss attributed to SIF activity.[4]

Understanding the cybercriminal and their ecosystem

Before diving into the crime, it's important to understand the cybercriminal.

Two traits cybercriminals share are **apathy** and **arrogance**. They are apathetic about committing crime or harming individuals. They are just as willing to steal money from a bank as they are to steal a senior citizen's social security benefits. Cybercriminals are also arrogant; they believe they can succeed against formidable opponents, such as a billion-dollar bank or multinational security company.

However, even the most apathetic and arrogant cybercriminal will fail without the necessary **skills**. The knowledge and skills to plan a cybercrime, gather the necessary data, execute, and cash out at the most opportune moment are unlike any required for traditional financial crime. These specialized skills vary dramatically from individual to individual.

As a result, most cybercriminals rely on a broad and growing ecosystem of forums, marketplaces, and dark web services to connect them with other cybercriminals who can augment their own capabilities. Only when they convene the right ecosystem can cybercriminals gain:

- Access to tools, services and Personally Identifiable Information (PII) (like social security numbers) required to perpetrate the cybercrime.
- Insight into the specific vulnerabilities and inefficiencies of the organization's ecosystem.
- Deep understanding and countermeasures to bypass security of the target organization.
- Visibility into the third-party services on which the organization relies.
- Tools and documentation (like fake driver's license) necessary to exfiltrate the proceeds of the crime successfully and without leaving a trail for the authorities to track.

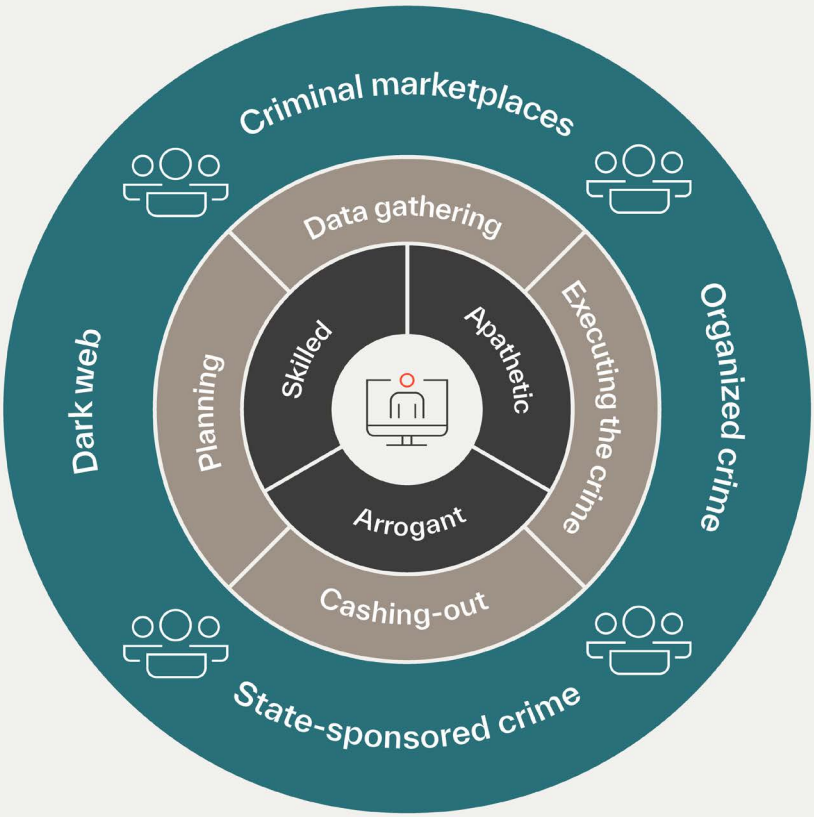

Why synthetic identity fraud succeeds

Synthetic identity fraud (SIF) succeeds due to a series of systemic failures that occur before the attacker even targets a specific enterprise. Understanding these failures is crucial for combating SIF.

Social Security Administration: U.S. Social Security Numbers (SSNs) were originally created using a specific pattern based on geographical location, making it easy to deduce a person's full SSN with just a few pieces of personal information. This led to a large and growing number of identity thefts. In 2011, the SSA began randomizing SSNs to combat traditional identity theft. The randomization of SSNs reduced a particular method of fraud but inadvertently enabled SIF by allowing criminals to fabricate SSNs using randomizer algorithms or exploit SSNs issued after 2011, like those of children or new citizens, for fraudulent activities.

Credit bureau failures: Credit bureaus do not know an individual exists until they apply for credit. So, for example, when a cybercriminal applies for credit using a new synthetic identity (SID) and the bank submits an inquiry to determine eligibility, the report will show a lack of credit history. But the inquiry itself triggers the credit bureau to create a record in its system for this entity. While the bank will most likely deny the application, the damage has been done at the credit bureau. A new entity has been created for a SID criminal to build a good credit history.

Credit piggybacking: Credit piggybacking is a strategy where an individual with a low credit score becomes an authorized user on someone else's credit account to benefit from the primary account holder's positive credit history. While legitimate credit piggybacking can help build a person's credit, criminals exploit this method to enhance the credit scores of their SIDs.

Here's how it works: A synthetic identity (SID) is added as an authorized user of a legitimate credit line account with a good credit history. This can occur by covertly convincing someone to add them, gaining access to their account, and adding trade lines. The fraudster usually does not intend to make purchases on the parent account. Once an authorized user is added, if the debt ratio is favorable, the available balance is high, and the card has been active long enough, the card's history becomes part of the fraudster's credit history during the ensuing credit reporting cycle. This can boost the credit score of a SID from 300 (typical starting credit score) to over 760 in 30–45 days. A skilled cybercriminal adds two trade lines to a single account and simultaneously boosts the credit rating of multiple SIDs.

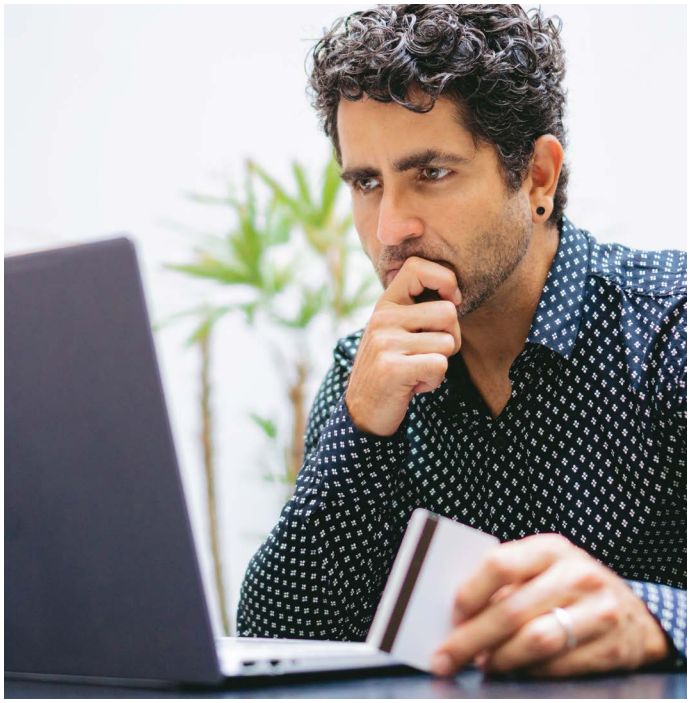

Traditional fraud models: Financial institutions' traditional fraud models often do not detect SIF. Victims—which often include children, the unhoused, elderly people, and incarcerated individuals—may remain unaware of the scam for years, making it challenging to create effective detection models. If a criminal creates a SID based on a 4-year-old child's SSN, that criminal may have a 14-year head start before the victim becomes aware of the crime. This is why traditional models are often ineffective; they are typically transactional in practice, looking at an applicant at a single point in time, not over a series of months or years.

Inconsistent definitions and reporting: Since SIF often occurs without detection many cases of loss get misidentified as poor underwriting decisions. Further, the definition of SIF varies across institutions and agencies, with some calling it identity theft, others cybercrime, and others synthetic identity crime. This makes it difficult for companies to share information about these crimes. The categorization of identity fraud, broadly, also varies. This makes the problem difficult to quantify. Also, no centralized reporting system exists, making it easier for criminals to operate across multiple institutions and coordinate their activities. Further, guidelines for categorizing and reporting SIF have not been established.

Technology innovation and automation: Financial institutions' increased reliance on AI/ML during the approval process has inadvertently created valuable data that criminals use to stay below risk thresholds in legacy fraud detection systems and make better SIF profiles to evade detection.

How synthetic identity fraud works

Synthetic identity fraud (SIF) is a methodical, long-game deception that exploits gaps in financial systems to extract money illegally. By strategically navigating credit reporting processes, fraudsters build credibility over time before cashing out. The following breaks down step-by-step how many of these bad actors establish, build, and ultimately profit from their SIDs.

Step 01: Creating the identity

Criminals often purchase personal identification information of children from the dark web. For approximately $2, you can buy a child's SSN, name, birth date, mother's maiden name, and address (place of birth). Initially, the criminal only needs the SSN. They combine the real SSN with a fabricated adult's name, date of birth, address, and phone number to create a synthetic identity.

Step 02: Applying for credit

Credit bureaus don't know you exist until you inform them. Cybercriminals simultaneously aim to create a credit report within all three credit bureaus (TransUnion, Experian, Equifax) by submitting a credit application that pings all three at the same time. This is called "tri-merging." This typically involves applications for car loans or mortgage refinances, regardless of whether the synthetic identity has a house or is trying to purchase an automobile. When the criminal submits the application, it will likely be denied. Still, since the credit bureau system hasn't seen that applicant before, it will create a record for the synthetic identity.

Step 03: Building the credit profile

Broadly speaking, the goal is to inflate the credit score as much as possible to cash out quickly. The first step involves using public data services to give the SID the appearance of legitimacy. When a credit pull occurs, creditors look for information that links you to an address, often using public databases. Fraudsters exploit this by ensuring web crawlers and databases associate the synthetic identity with an address, email and phone number. Meanwhile, the fraudster opens reward cards and may create social media accounts to legitimize the false identity further. Next, they might use secured credit cards, which don't boost the credit score but provide immediate financial gain since some secured credit card companies offer more credit than the deposit amount, allowing fraudsters to profit and fund other activities.

Step 04: Cashing out

Once the synthetic identity has a high credit score, criminals can cash out by obtaining loans, credit cards, furniture, cars or other valuable products. The more skilled criminal may try to grow and obscure the crime by making payments for some time and then cashing out to extend the profit timeline while reducing the likelihood that that identity will be flagged as fraud. The act of cashing out a synthetic identity crime will involve an act of money laundering because it involves the transfer of illicit funds into a different account to conceal its origins. To accomplish this, fraudsters often need to open bank accounts, which typically require a driver's license. It is yet another example where the ecosystem supports the cybercriminal, where nearly perfect fake driver's license can be acquired for $40 to $80.

Why traditional detection fails

Current customer onboarding systems are primarily designed to evaluate an individual or business at the point of application for credit or loans. These systems attempt to verify the applicant's documents, assess their creditworthiness, and review their financial profile. The process includes checking credit scores, payment history and existing debts. Underwriters then analyze the requested credit limit or loan amount and the applicant's ability to repay. Unfortunately, these systems often fail to track the identity past the onboarding stage and throughout its lifecycle. Once an identity is onboarded and deemed low risk, it is not continuously monitored with the same rigor. This gap allows SIDs to evolve and become high-risk without detection.

Financial institutions need to adopt more dynamic and integrated approaches to identity verification and monitoring to address these challenges. This includes using advanced technologies like AI/ML and Graph Analytics to continuously analyze and flag suspicious activities throughout an identity's lifecycle.

How synthetic identities go undetected

- **Static verification:** Onboarding systems typically perform a one-time verification process. After the initial check, the identity is not subjected to continuous scrutiny, making it easier for SIDs to remain undetected as they build a credit history.
- **Lack of integrated monitoring:** The siloed nature of many organizations makes it challenging to observe all aspects of a customer's activities, making it more difficult to spot patterns of fraud. For example, many financial institutions use separate onboarding and customer management systems. These systems often do not communicate effectively, leading to a lack of comprehensive tracking of an identity's activities over time.
- **Evolving fraud tactics:** Fraudsters continuously adapt their methods, making it challenging for static systems to keep up. SIDs can start as low-risk and gradually engage in fraudulent activities, exploiting the lack of continuous monitoring.
- **Resource constraints:** Continuous monitoring of all accounts can be resource intensive. Financial institutions may prioritize high-risk accounts, leaving low-risk accounts with less scrutiny.

Steps to fight back

Key measures enterprises—and financial institutions, in particular—can implement to address the many challenges posed by SIF include:

- **Analyze current losses and exposure:** Banks should start by analyzing their current credit card or loan charge-off losses to determine their exposure to SIF crime. Developing categorization methods to track and report fraud loss and activity more effectively will provide a clearer understanding of the impact and help devise targeted countermeasures.
- **Understand the regulatory landscape:** Stay informed about privacy regulations, AI bias reporting requirements, and data collection and destruction mandates. Compliance with these regulations enables banks to implement anti-fraud measures without violating legal protections.
- **Use encryption and tokenization:** Implement encryption and tokenization tools to mask personal identification information while retaining its fidelity. This supports the creation of detailed relationship graphs, which are vital for identifying SIDs without compromising data privacy and security standards.
- **Establish a rapid response team:** Create a rapid response that's equipped to develop and implement countermeasures swiftly, ensuring that the bank can respond to emerging threats quickly.
- **Use large, diverse, and accurate datasets:** Access to a large, diverse, and accurate dataset is fundamental for effective onboarding and transaction monitoring. A comprehensive dataset enhances an enterprise's ability to train models and detect anomalies indicative of SIDs. Advanced data management and pipeline solutions capable of rapidly integrating diverse datasets are essential. These platforms allow for the seamless combination of various data sources and formats, providing a holistic view of customer and merchant activity, enabling more accurate anomaly and fraud detection.
- **Implement a flexible rules engine platform:** A flexible rules engine platform allows banks to adapt quickly to new fraud patterns. By continuously updating and refining rules based on the latest intelligence, banks can stay ahead of fraudsters and reduce the incidence of SIF.

- **Employ advanced AI and machine learning (ML) models:** AI and ML tools and models can play a key role in helping to detect and mitigate SIF. For example:
  - Auto-encoders, which are neural networks used for unsupervised learning, help clean data by learning to compress and then reconstruct it, effectively filtering out noise.
  - Gradient boosting is a technique that combines multiple weak learners to create a strong predictive model. A weak learner is a model that performs slightly better than random guessing, often having limited predictive power on its own. Gradient boosting is effective in identifying SIF patterns by sequentially adding weak learners, typically decision trees, each correcting the errors of the previous ones. This iterative process enhances the overall model's accuracy and robustness.
  - Support Vector Machines (SVM) are powerful for classification tasks, sorting data into different categories, and are particularly suitable for complex datasets with many features, making them adept at detecting SIF.
  - Random Forest, a robust machine learning model, analyzes multiple decision trees and averages their results to improve accuracy, handling large datasets and providing feature importance scores to aid fraud analysts.
  - Principal Component Analysis (PCA) reduces noise by focusing on the most essential features of the data.
  - Generative AI models can evaluate unstructured data, such as social media, and generate synthetic test and training data that mimic real-world SIF scenarios.
  - A federated learning approach, like NVIDIA Flare, enables multiple financial institutions to collaboratively train a fraud detection model without sharing sensitive customer data directly, improving accuracy while maintaining data privacy.
- **Leverage graph analytics:** Advanced graph analytics tools, using nodes (entities) and edges (relationships) to model data, are invaluable for visualizing and analyzing complex relationships within data and uncovering hidden patterns and relationships among transactions.
- **Curate third-party services and tools:** Understanding and curating third-party services, tools, data, and software can significantly augment and enhance your organization's SIF strategy. This includes:
  - **Third-party data providers:** Partner with reputable data providers to access additional data sources that can enrich your existing datasets. This can include credit bureaus, public records, and other external data services.
  - **Fraud detection software:** Invest in advanced fraud detection software that leverages machine learning and artificial intelligence to identify suspicious activities and potential SIDs. These platforms often train models on data from their entire client base, which can enhance their accuracy. These solutions usually offer alerting services among participating clients to share intelligence in real-time.
  - **Identity verification services:** Use third-party identity verification services that offer biometric authentication, document verification, and other advanced techniques to ensure the authenticity of customer identities.
  - **API integrations:** Develop an application framework that supports seamless integration of third-party API services into your existing systems, enhancing your overall fraud detection and prevention capabilities.

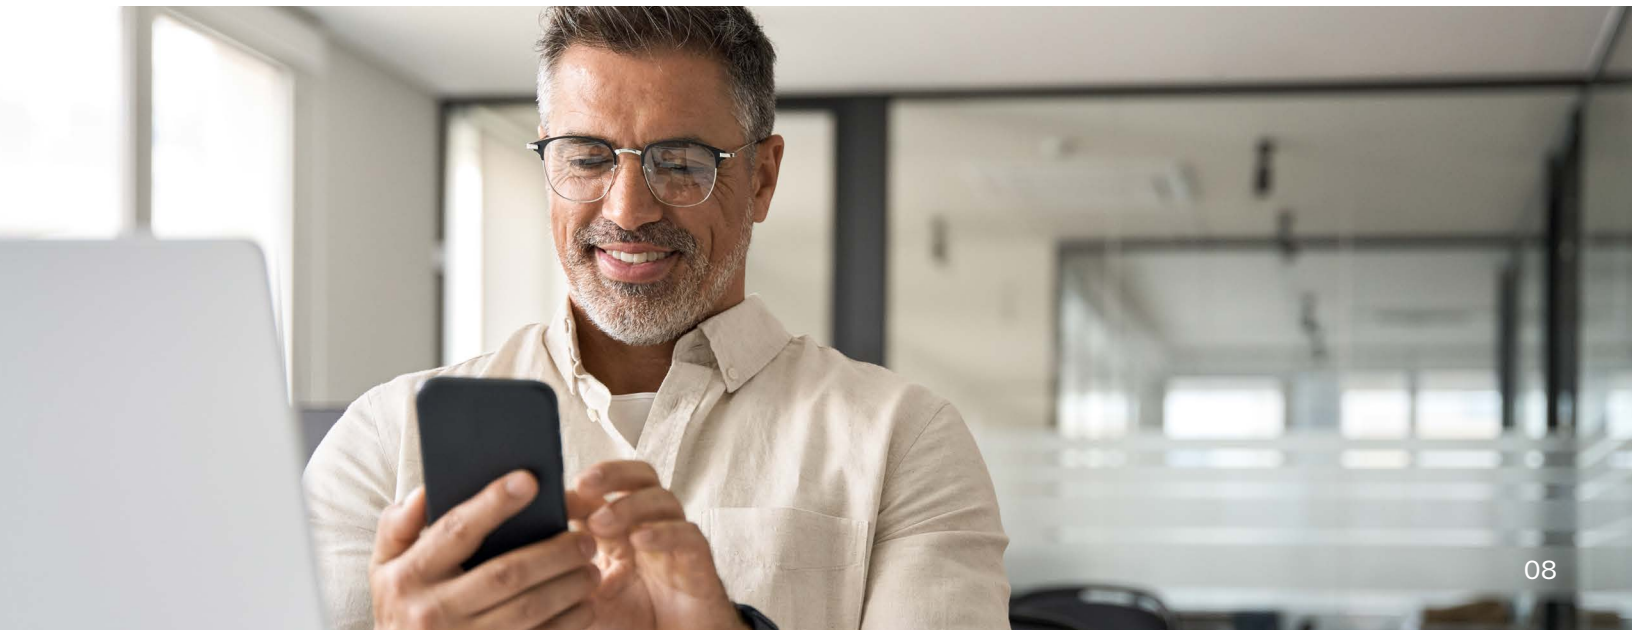

Kyndryl's perspective

Synthetic identity fraud (SIF) is a complex and rapidly growing threat that requires a comprehensive understanding of the criminal, the ecosystem, and the systemic failures that enable it. Working with a trusted advisor to navigate the many complexities of the field is critical for developing effective strategies to combat this pervasive form of financial crime. It is possible to significantly enhance the ability to spot and thwart these crimes.

Still, any solution must keep pace with the fast-evolving and innovative criminal ecosystem. Firms must be willing to adapt, break down data silos, and partner with industry regulators and a growing number of innovative consultants and solutions providers who specialize in advanced fraud detection and cybersecurity.

That work starts by gaining access to valuable insights into the latest trends and best practices in defending your organization against the SIF threat, and the organized criminal enterprises that enable the practice. Enterprises should seek expert guidance on identifying vulnerabilities and developing robust detection and prevention strategies. Tailored managed services offer ongoing monitoring and support, allowing firms to stay agile and responsive to new threats as they emerge.

Author

Robert Wallos
robert.wallos@kyndryl.com

Robert Wallos is the Director of Enterprise Architecture at Kyndryl, with extensive experience in banking, payments, and capital markets. He is a recognized thought leader in the financial services industry with expertise spanning cloud computing, AI, fraud mitigation, open banking, automation, and ultra-low latency pub/sub architectures.

Sources

[1] Global Compliance Institute. July 2024.
[2] Deloitte Center for Financial Services. July 2023.
[3] U.S. Federal Reserve. July 2019.
[4] Global Compliance Index. July 2024.
[5] Identity Theft Resource Center. January 2024.

kyndryl