iProov

Stolen to Synthetic:

The Evolution of Identity Fraud and the Need for Resilient Identity Verification

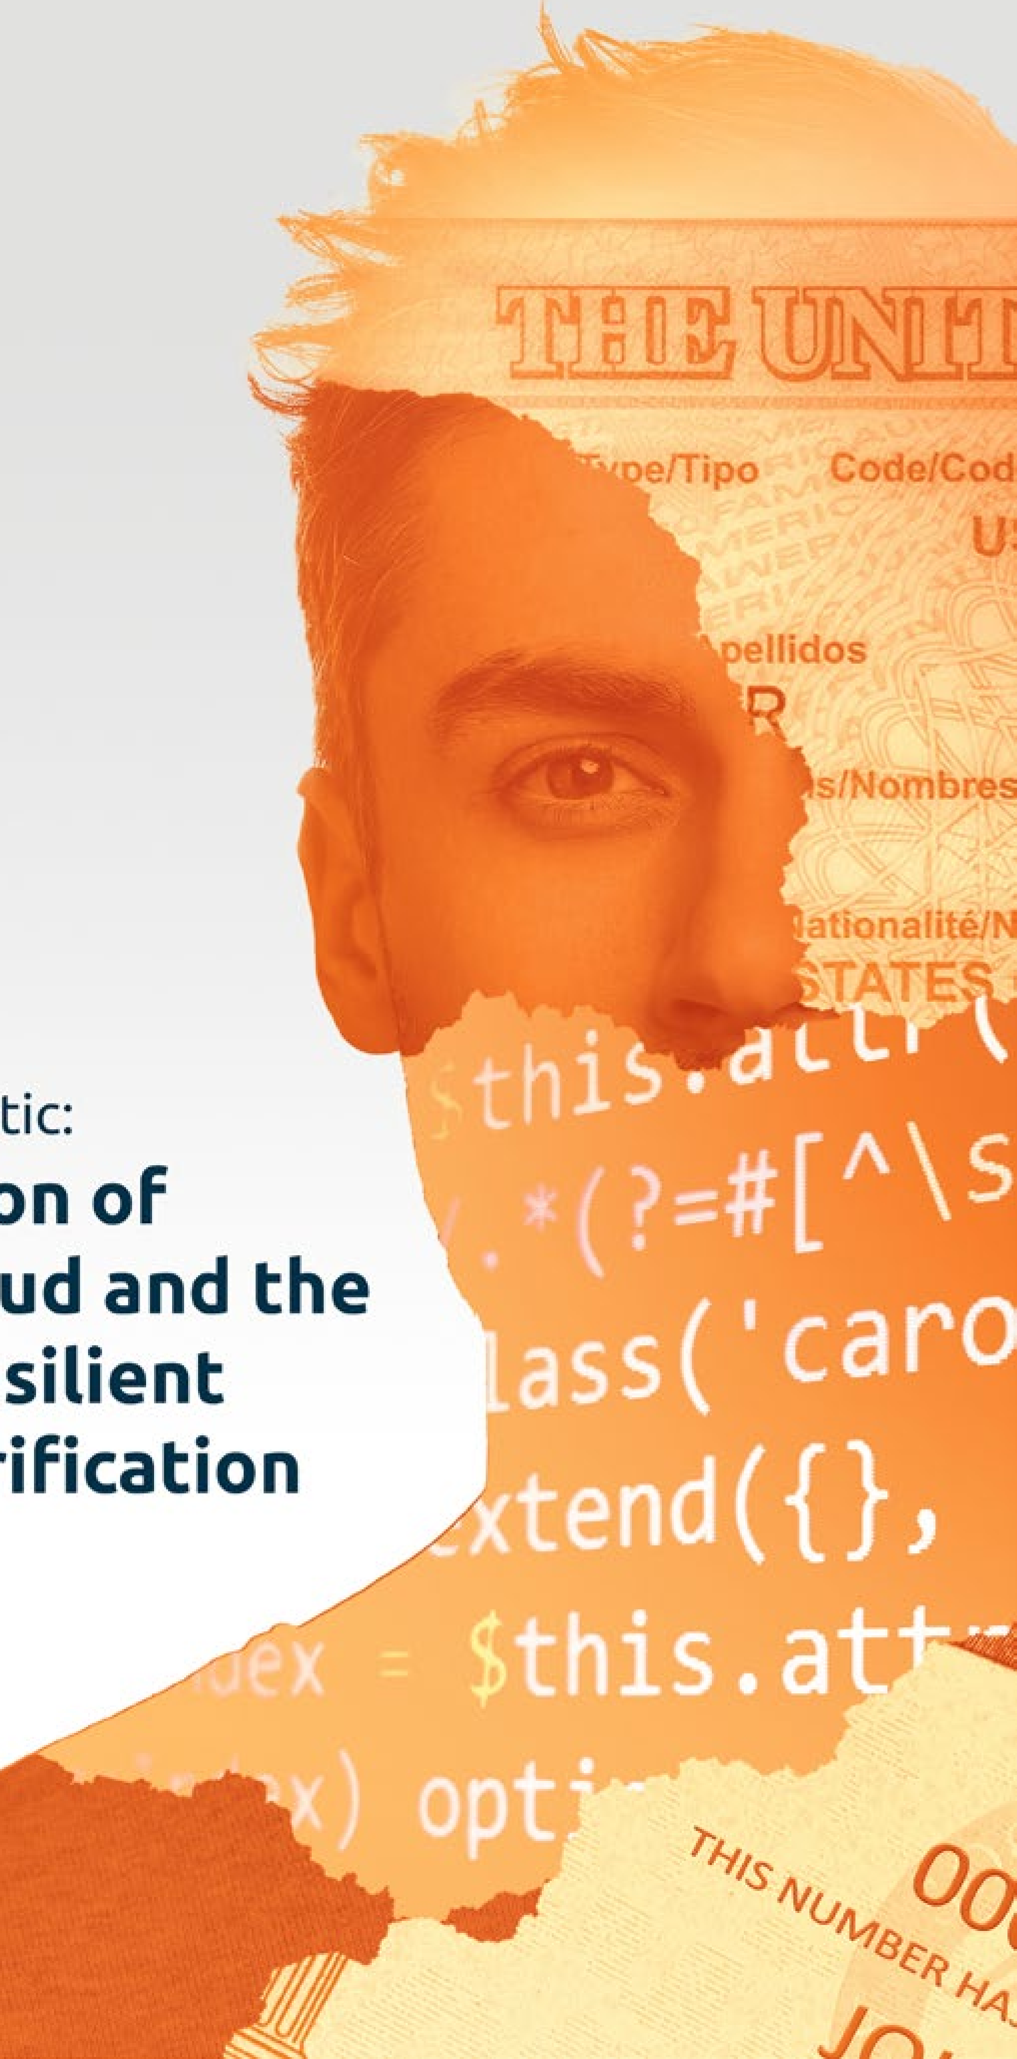

Introduction

In 2020, Adam Arena and co-conspirators were indicted on charges of attempting to steal over $1m from various banks.[1] The fraudsters developed a complex web of synthetic identities and used them to obtain loans and credit cards from financial institutions that they never intended to repay. They had even created shell companies to further legitimize the synthetics and apply for larger loans.

The scheme was so successful that months later, Arena did it again.[2] This time, his target was the US Government's Paycheck Protection Program, a COVID-19 emergency relief program designed to help businesses cover payroll and other costs.

Synthetic Identity Fraud, or 'SIF', is 'the use of a combination of personally identifiable information (PII) to fabricate a person or entity to commit a dishonest act for personal or financial gain.'[3]

McKinsey & Company estimates synthetic identity fraud as the fastest-growing financial crime in the United States[4]. Successful SIF schemes, while highly lucrative for criminals, are in no way victimless. They come at the detriment of organizations, governments, and society at large.

Beyond defrauding financial institutions and public sector programs, SIF has wider societal ramifications. It facilitates money laundering, human trafficking, and terrorist funding. The US Financial Crimes Enforcement Network (FinCEN) has named SIF as a national security priority,[5] and the Federal Reserve recently launched its Synthetic Fraud Identity Mitigation Toolkit to help financial services organizations combat SIF.[6]

To stem this threat, organizations must deploy secure identity verification at onboarding. This report examines how identity verification and fraud detection must progress to stay ahead of adversaries. It illustrates how biometric face verification can thwart synthetic identity fraud before it can cause damage.

Finally, it looks at how fraudsters have harnessed generative AI to leverage synthetic identities with synthetic imagery and how biometric systems can provide resilience against this evolving threat.

[1] Frankenstein Fraud: How Synthetic Identities Became the Fastest-Growing Fraud Trend, ASIS International, 2021
[2] Man Gets 4-12 Years In Prison In ID Fraud Scheme, Patch, 2023
[3] Synthetic Identity Fraud Mitigation Toolkit, The Federal Reserve, 2023
[4] Fighting Back Against Synthetic Identity Fraud, McKinsey & Company, 2019
[5] Prepared Remarks of FinCEN Acting Deputy Director Jimmy Kirby During the 2022 Federal Identity Forum & Exposition (FedID), Financial Crimes Enforcement Network, 2022
[6] Ibid

Synthetic Identity Fraud in Numbers

20bn

Business losses due to SIF in 2020
Federal Reserve Toolkit

$97,000

The average amount lost per SIF case
PYMNTS

85%

The percentage of synthetics not detected by traditional fraud models
LexisNexis

$4

The average price for a stolen SSN on the dark web
Atlas VPN

1m

The number of synthetic identities found in the credit industry in a 12-month period
Equifax

The Evolution of Identity Fraud

Identity fraud is rapidly evolving. In the past, it was often characterized by one individual stealing another's PII to impersonate them for illicit financial gain. This is considered third-party fraud – or identity theft. Fraudsters have since graduated to the more insidious threat of SIF.

SIF is now the leading threat to organizations as it has overtaken identity theft in frequency and complexity, accounting for 80-85% of all identity fraud cases[7]. Governments and organizations globally are feeling its impact. According to one study, 46% of organizations were targets of synthetic identity fraud in 2022[8].

What's Changed?

The Explosion of the Digital Ecosystem

The Global Pandemic accelerated the pivot from previously in-person services to entirely remote. From making payments to enrolling in a public benefits program, organizations are looking to meet the user demand for low-friction, end-to-end digital experiences.

The challenge here is ascertaining that a remote user is who they claim to be, or even a real person at all, with a high degree of confidence.

Availability of Personal Identifiable Data (PII)

In third-party fraud cases, the fraudster is linked to the target (family or acquaintance) 99.1% of the time.[9] This is not the case for SIF. PII is readily available to purchase from the dark web or attainable for free from social media.

Wide-scale data leaks, meanwhile, have exposed the PII of millions to nefarious actors. In 2021, a zero-day attack against Facebook led to the leakage of 533 million users' personal details and contact information[10]. The abundance of PII enables criminals to create synthetic identities using information from people they've never met.

Rise of Generative AI

Sophisticated generative AI tools are no longer a thing of the future. Now, low-skilled criminals can create advanced, highly-sophisticated synthetic imagery that make synthetic identities seem all more real.

We discuss this further on **page 11**.

[7] The Changing Face of Identity Theft, Federal Trade Commission, 2022
[8] The True Numbers Behind Deepfake Fraud, Help Net Security, 2023
[9] Uncovering Synthetic Identity Fraud, LexisNexis, 2021
[10] 533 million Facebook users' phone numbers and personal data have been leaked online, Insider, 2021

Reaching the Unbanked

New-to-market consumers, including young people, recent immigrants, and historically unbanked populations, represent a growth market for financial institutions. In the US, for example, the underbanked spend $189bn annually on alternative financial services, like payday loans.[11] In a race to acquire these people, financial institutions must open their doors to those with little to no credit and, at times, no financial history.

Yet, these individuals also provide the ideal cover for synthetic fraudsters, who pose as the unbanked or underbanked to open new credit and bank accounts.

[11] Investing In The Underbanked: An Opportunity At The Intersection Of Fintech And Proptech, Forbes, 2023

How to Create a Synthetic Identity

The Bare Necessities

Fraudster + Fake name + Social Security or ID Number = Synthetic Identity

How to Further Validate a Synthetic:

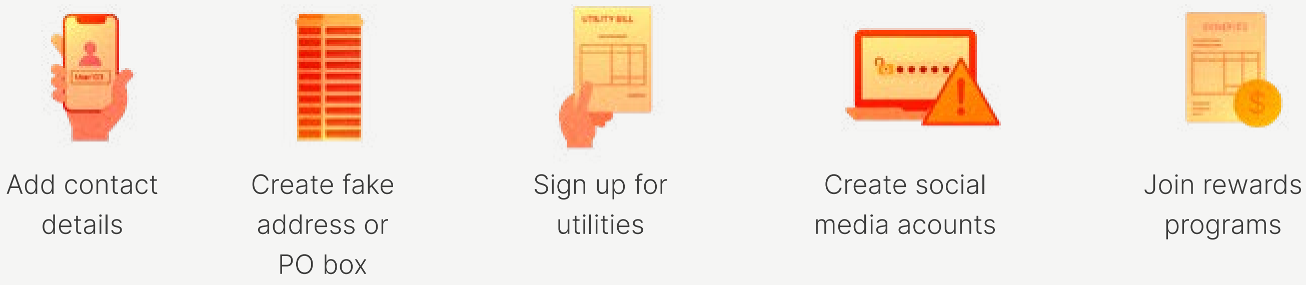
Add contact details

Create fake address or PO box

Sign up for utilities

Create social media acounts

Join rewards programs

The Scalability of SIF

Although synthetic fraudsters also target public services and banks, it is in the credit sector where most synthetic identities are flagged.[12]

This is due to scalability. Fraudsters can 'piggyback' by adding a synthetic onto the credit account of a genuine customer. The synthetic benefits from the real person's credit history. Once the synthetic has a good credit history and can start applying for credit of their own, the fraudster creates more synthetics, which 'piggyback' off the first synthetic.

Eventually, the fraudster has a web of synthetics with their own credit lines. They can nurture these for years before eventually 'busting out.'

Furthermore, fraudsters collaborate to gain the most value from their SIF schemes. They work together to create multiple synthetic identities, increase payouts, and avoid detection.

[12] Synthetic Identity Fraud: A Look Behind the Mask, Equifax, 2019

Use Case: Using SIF to Defraud a Credit Institution

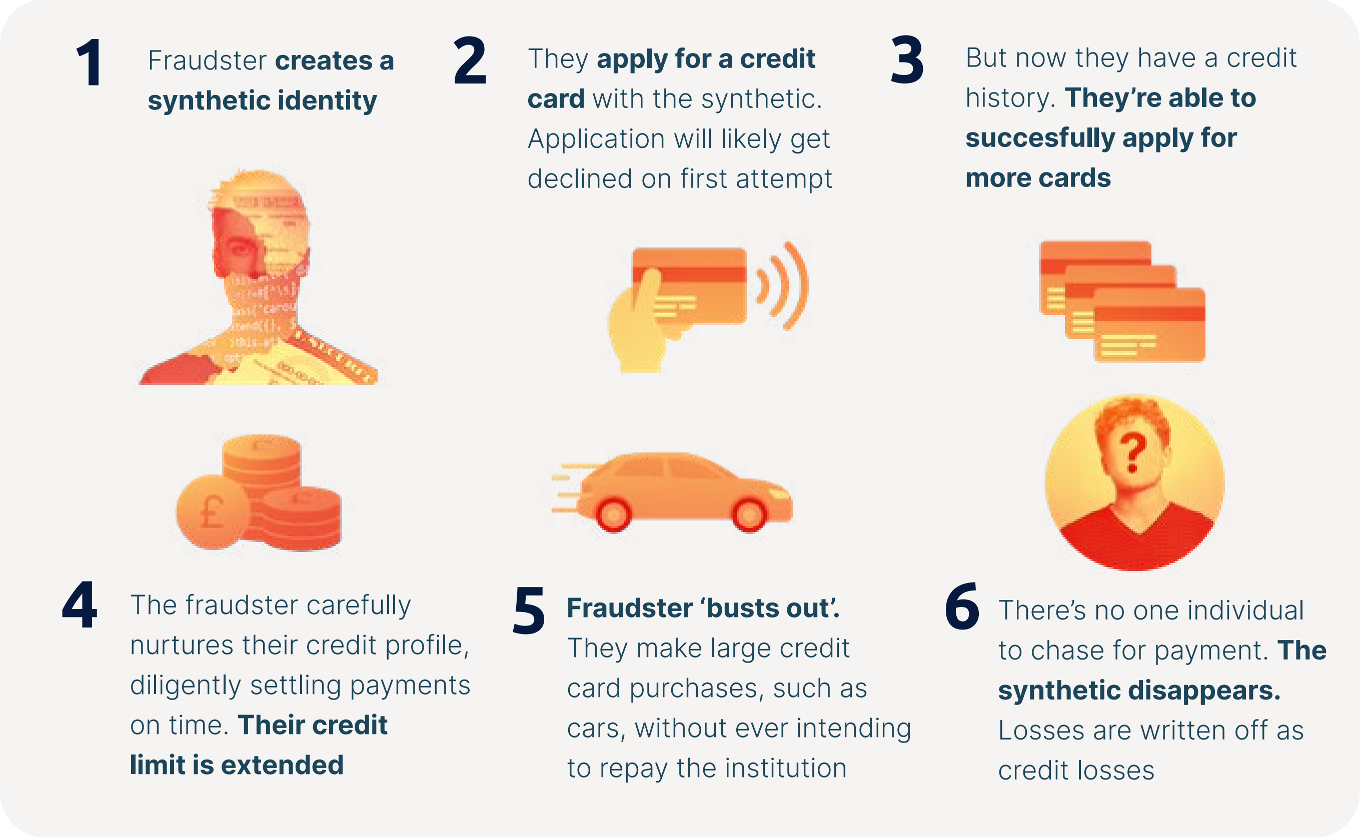

1 Fraudster **creates a synthetic identity**

2 They **apply for a credit card** with the synthetic. Application will likely get declined on first attempt

3 But now they have a credit history. **They're able to succesfully apply for more cards**

4 The fraudster carefully nurtures their credit profile, diligently settling payments on time. **Their credit limit is extended**

5 **Fraudster 'busts out'.** They make large credit card purchases, such as cars, without ever intending to repay the institution

6 There's no one individual to chase for payment. **The synthetic disappears.** Losses are written off as credit losses

The Issue With Traditional Fraud Detection

Traditional fraud detection models are designed to detect third-party fraud. They look at the activity, not the individual themselves. If a known genuine user starts behaving suspiciously, such as making large purchases in quick succession or logging in from an unknown location, the detection models will be alerted.

> “Fraud detection often focuses primarily on preventing unauthorized account access and transactions by an external party – and doesn’t always look for a party who is not real.”
> – **Federal Reserve, Synthetic Identity Fraud Mitigation Toolkit**

Here Are Several Ways That Synthetic ID Fraudsters Evade Detection

Synthetics Behave Like Real People

With third-party fraud cases, the perpetrator's actions arouse suspicion. They make large payments to a new recipient, for example.

This is not the case with synthetic identity fraud. It is not uncommon for people to have poor credit or default on loans. Synthetics can exist for years without doing anything that alerts the authorities. This makes it incredibly challenging to detect SIF in the user lifecycle.

Lax Onboarding Processes

As the threat landscape evolves, so must the technology and processes needed to verify that the remote individual is who they say they are.

Unfortunately, organizations typically rely on source documents at onboarding that can be fabricated, such as social security numbers or identity documents.

The Synthetic Can Re-authenticate With the Same Credentials

Once a synthetic has created an account with an organization, they can return to the service using the same credentials, whether email, password, or biometric. As the credentials match the identity of those onboarded in the first instance, the fraudster can use them again without arousing suspicion. They're then free to continue to use the service for malicious purposes.

Detecting SIF With Biometric Face Verification at Onboarding

As synthetics behave like real people, it's incredibly challenging to detect them in the user lifecycle. According to the Federal Reserve, "more security is needed at onboarding" to detect if an unknown individual is a real identity before they become an established user.

To achieve this, many organizations have adopted biometric face verification for onboarding users. A critical element in biometric technology is liveness detection, which is able to determine whether a remote user is a real, 'live' person and not a non-living spoof. This has many advantages over other verification methods:

Assurance: The security of biometric face verification does not rely on a face being a secret but rather that it is an inherent factor. Unlike many verification technologies, it does not rely on something you have like a token or something you know like a password. For synthetic ID fraudsters, forging a face is exponentially more challenging than forging other credentials, like a social security number. But not impossible, as we illustrate next.

User convenience: People have their faces wherever they go, meaning they can use it to onboard or authenticate wherever they are. Unlike other credentials, faces never have to be reset, and can't be forgotten at home, mitigating unnecessary friction for the user.

No onus on the user: The 2023 Whitehouse Cybersecurity Strategy states that end-users should not bear the burden of their safety when using online services.[13] Instead, "it must be the responsibility of the owners and operators of the systems." If deployed correctly, biometric face verification doesn't burden the user. They have nothing to remember or forget, nor need to speak with an operator.

Biometric Face Verification For Onboarding

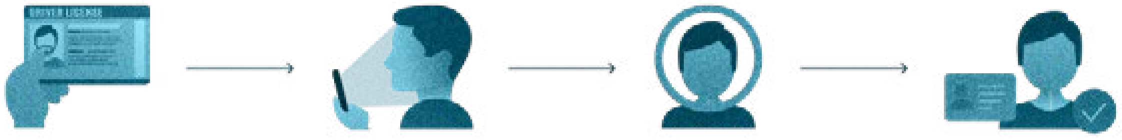

An individual captures information from their trusted government-issued ID

They take a selfie. Face matching determines if the two images are the same

The system checks whether the image capture of the individual is 'live' and not a spoof

Biometric face verification check complete

[13] National Cybersecurity Strategy, The White House, 2023

How Fraudsters Harness Generative AI to Undermine Face Biometrics

Thanks to the widespread popularity of tools like ChatGPT, generative AI has launched to the forefront of the technological zeitgeist. While the technology has many positive uses, criminals have also harnessed it for illicit purposes.

Generative Adversarial Networks (GANs) are a branch of generative AI that use deep learning methods, such as convolutional neural networks.[14] Two systems, a generator, and a discriminator, operate in an adversarial zero-sum game. The generator creates imagery, and the discriminator deems it real or not. Eventually, the generator creates imagery that is realistic enough for the discriminator to deem as real.

GANs can be used to create synthetic imagery of people who do not exist. Their applications are multifaceted. They have even been used to create synthetic fashion models, thereby removing the cost of the model, photographer, and set.

Beyond putting people in the fashion industry out of business, GANs are also used to leverage synthetic identity fraud schemes. As illustrated below, GANs create synthetic imagery to match the false identities, enabling fraudsters to bypass biometric face verification technology deployed at onboarding.

Generative AI has become increasingly accessible in recent years. Available as toolkits from code depository domains, even coming with accompanying tutorials, this technology has enabled low-skilled criminals to launch advanced biometric attacks.

In 2022, iProov witnessed the emergence and growth of face swaps, sophisticated synthetic imagery that combines the traits of one face with the motions of another.[15] A 295% rise in face swaps from H1 to H2 indicates the increasing ability of threat actors to create and launch sophisticated synthetic imagery.

[14] A Gentle Introduction to Generative Adversarial Networks (GANs), Machine Learning Mastery, 2019

[15] iProov Biometric Threat Intelligence Report, iProov, 2023

How Fraudsters Harness Generative AI to Undermine Face Biometrics

1 Fraudster **creates a synthetic identity**

2 Using generative AI, they create imagery of **a person who doesn't exist**

3 The fraudster doctors an **ID to show the new face**

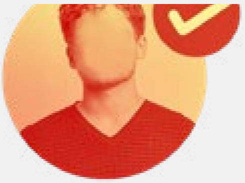

4 By matching the synthetic face with the fake ID, **they are able to circumvent face matching technology**

5 Fraudster is able to create a **new account** with a financial institution or government agency

6 The fraudster **uses the account for criminal purposes,** like credit fraud or money laundering

Launching Synthetic Imagery Attacks

Once synthetic imagery is created from generative AI, the fraudsters must deploy it against face biometric technology. There are two ways to do this:

Presentation attack: Present the synthetic imagery on a device and show it to the camera
Digital injection attack: Digitally inject the synthetic imagery directly into the data stream.

The biometric threat landscape is rapidly evolving. Presentation attacks and Presentation Attack Detection (PAD) are now well understood. Many biometric vendors have been accredited for detecting this threat.

Yet, digital injection attacks are now the far greater threat in terms of sophistication, frequency, and scalability. Biometric face verification with PAD alone cannot defend against this threat, leaving a vulnerability gap that synthetic ID fraudsters can exploit. Organizations must provide resilience to the most advanced threats to thwart SIF.

5X More frequent

Digital injection attack vs. presentation attacks in 2022

iProov Biometric Threat Intelligence Report 2023

	Presentation Attacks	Digital Injection Attacks
Deployment	An artifact is presented to the camera	Imagery is injected directly into the video stream
Detection	Detected via clues in the imagery	Detected either via analyzing metadata or imagery-based testing
Scalability	Limited in scale. Threat actors can deploy one attack at a time	Unlimited in scale. Threat actors can create highly automated attack machines
Testing	Existing globally recognized standards for Presentation Attack Detection (ISO/IEC 30107)	No existing globally recognized standards for Digital Injection Attack Detection

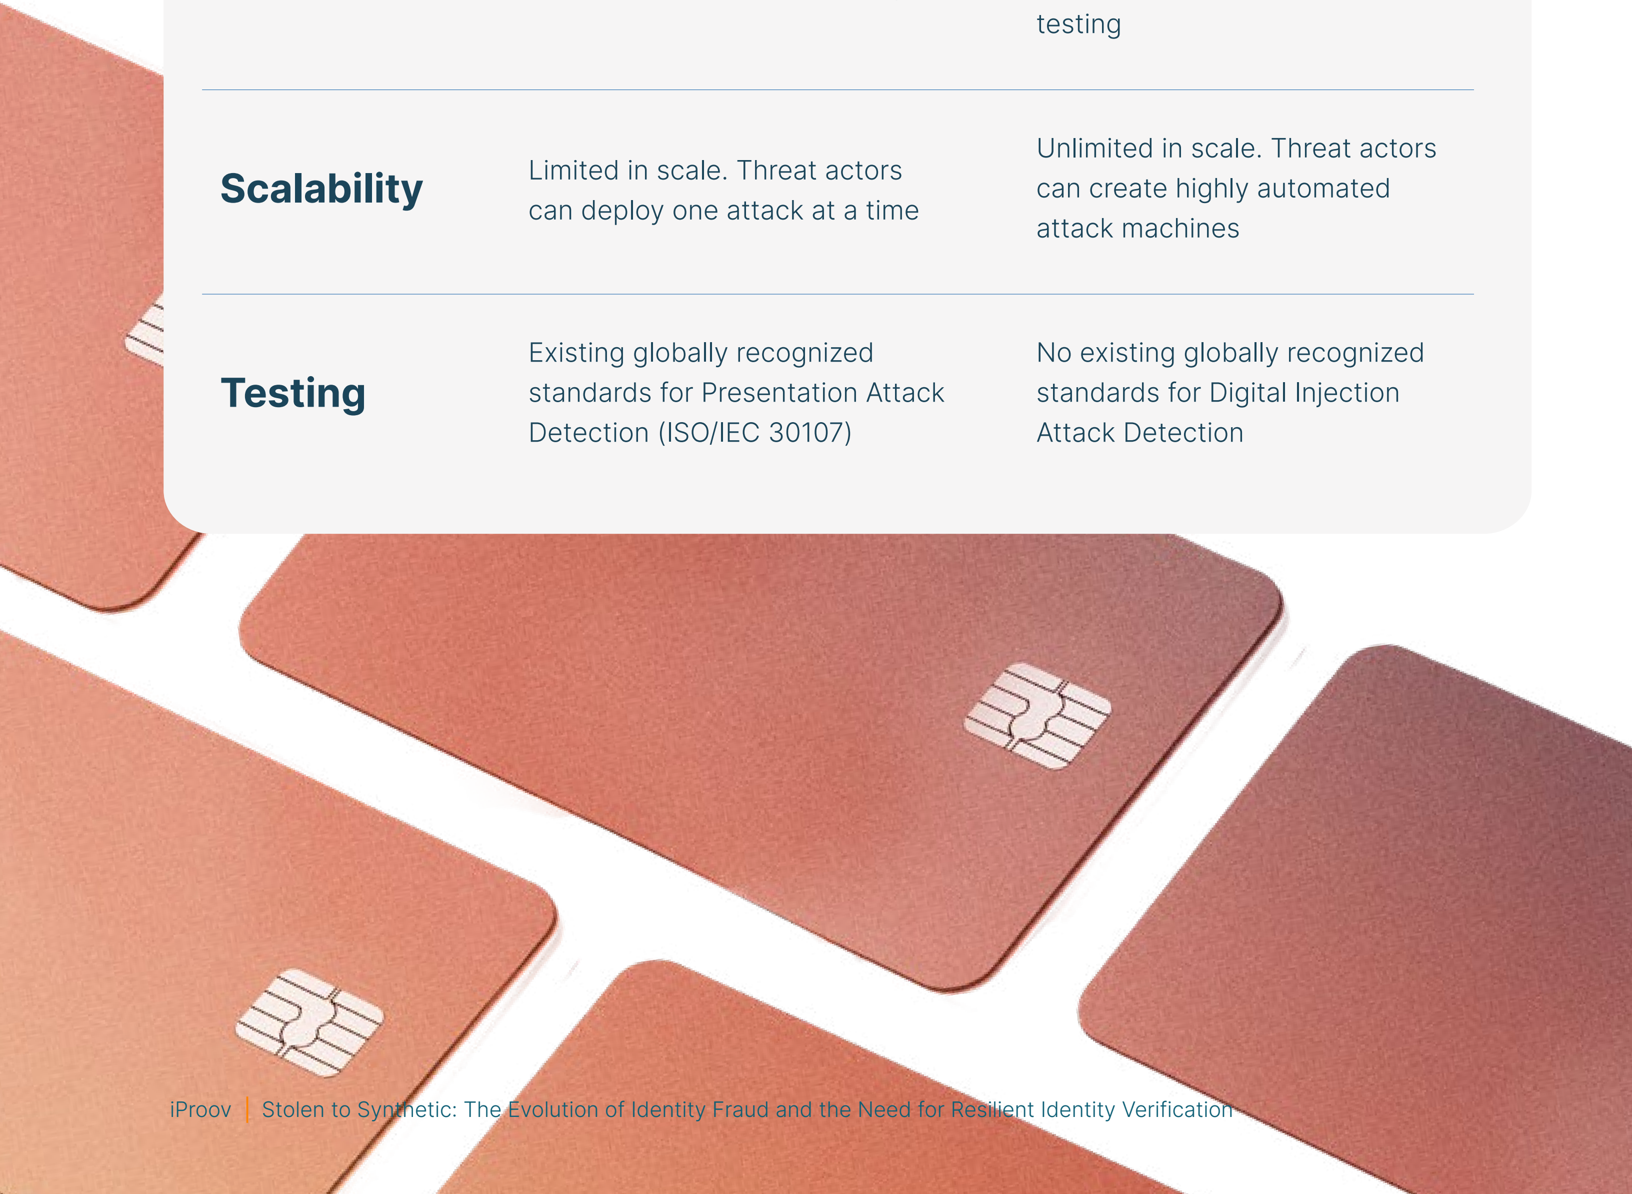

Providing Resilient Biometric Verification at Onboarding

To successfully mitigate the threat of SIF, organizations must deploy biometric verification at onboarding. But not all biometric face verification is created equal. To secure high-risk use cases, such as onboarding with financial institutions or online government programs, the biometric face verification solution deployed must be resilient against the biometric threat landscape. To achieve this, organizations and the biometric vendors they employ must deploy one-time biometrics.
A one-time biometric ensures that a remote user is not just 'live' at the point of onboarding but also verifying in real-time**.**

Active one-time biometrics:

Not Inclusive: They request the user to turn their head in different directions a number of times or reading out loud a set of characters that changes each time somebody verifies their identity. They can impair accessibility and inclusivity. Those with physical or cognitive disabilities may be unable to perform certain actions or follow complex instructions.
Security Compromise: Active solutions also present a security issue. The authentication process – i.e., the sequence of actions the user is asked to perform – is exhibited to threat actors, providing them with the information they need to reverse-engineer the technology.

Passive one-time biometrics:

Inclusive: The user does not perform actions. They look at the camera and the authentication process is complete.
Secure: With passive one-time biometrics, the challenge-response mechanism is randomized, making the authentication process unpredictable, impervious to replay attacks and highly challenging to reverse-engineer.

Passive one-time biometric user journey

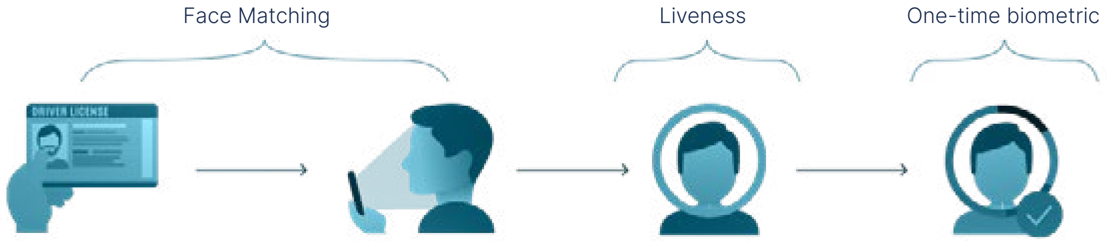

An individual captures information from their trusted government-issued ID

They take a selfie. Face matching determines if the two images are the same

The system checks whether the image capture of the individual is 'live' and not a spoof

Uses multi-dimensional technology, with random challenge response to check that the individual is not a generative AI or other sophisticated synthetic media

Summary

Synthetic ID fraudsters play the long game. They may lie in wait for years, building the legitimacy of their identities or diligently tending to their credit profile. When they strike and disappear, it's all too late, the chance of mitigation has evaporated.

As such, it has become increasingly difficult for governments and organizations to ascertain which of their users are real and which are fictional. This makes SIF not just a fraud issue but also a compliance problem. Allowing synthetic identities access to online services fundamentally undermines Know-Your-Customer (KYC) processes.

What's more, the fuel that powers SIF is not going away. In 2022, 1,774 organizational data compromises exposed the PII of over 392 million individuals globally.[16] This personal identifiable information, obtained through cybercriminal activity, melded with generative AI tools, creates sophisticated synthetic identities that are becoming all the more believable.

Organizations must act now to stem this threat. Robust identity verification at onboarding is required. Biometric face verification remains the most secure and usable way to ensure a remote user is who they claim to be.

Yet, not all biometrics are equal. The threat landscape is evolving rapidly and has outpaced legacy Presentation Attack Detection systems. To combat synthetic identity fraud, organizations must also fight the current and future methods fraudsters use to legitimize and deploy SIF schemes. This means deploying digital injection attack detection to thwart synthetic imagery circumventing biometric systems. Deploying one-time biometrics at onboarding is the most effective way to incapacitate synthetic identity fraud before it can cause damage.

[16] Data breaches are increasing at a rapid speed. Here's what can be done, World Economic Forum, 2023

iProov is used by leading organizations worldwide to reduce the risk of identity fraud.
The Federal Reserve has listed iProov as a Synthetic Identity Fraud Mitigation Provider.

Financial services clients include UBS, ING, Rabobank, and Knab.

Government clients include the U.S. Department of Homeland Security,
the UK Home Office, the Australian Tax Office, and GovTech Singapore.

contact@iproov.com

iproov.com