

Fraud Research

Paid in Full

How card fraudsters in the GCC use government payment portals to cash out stolen cards

Table of contents

01 Disclaimer 04

02 Acknowledgements 06

Authors 06

03 Introduction 07

04 Key Findings 08

05 Attribution 10

Threat actor profiles 11
Scam profile 15
Victimology 16

06 Government Payment Cash-Out Scheme 17

From wallet top-ups to government portals 17
Stage 1: Acquisition and targeting 19
Stage 2: The phishing toolkit 20
Stage 3: Credential and authentication harvest 24
Stage 4: The checker layer (account takeover) 25
Stage 5: The cash-out (government portal and customer) 26
Stage 6: The underground market and customer recruitment 27
Stage 7: Laundering and exit rails 32
Why the scheme works 32
Detection in practice 32

07 Towards a predictive defence: Group-IB's cyber fraud fusion approach 34

Scenario: predicting and disrupting the bill-discount cash-out 35
Why fusion matters 36

08 Conclusion 36

09 Recommendations 37

For affected institutions (banks and issuers) 37
For detection and monitoring teams 37
For government and portal operators 37
For the ecosystem 37

10 MITRE ATT&CK 38

11 Fraud Matrix 39

12 Indicators Of Compromise (IOCs) 39

01 Disclaimer

01

This report is provided for informational purposes only. The findings, assessments, and recommendations are based on data available to Group-IB at the time of publication and are subject to change as new intelligence emerges.

02

While every effort has been made to ensure the accuracy of the information, Group-IB makes no warranties, express or implied, regarding completeness, reliability, or suitability for any particular purpose.

03

Assessments and attributions in this report carry the confidence levels expressed in the text. Where activity is attributed to accounts, channels, personas or groups, Group-IB makes no assertion as to the identity of any individual behind them beyond what is expressly stated. Names, handles and display names appearing in this report are aliases attributed to observed activity, and are not intended to refer to, or to imply any connection with, any real person bearing a similar name.

04

The institutions, billers, payment providers and platforms referenced in this report are described solely to document how threat actors abused, impersonated or routed activity through them. They are understood to be themselves subject to that abuse. No involvement, wrongdoing, negligence, liability or compliance failing on the part of any such institution, provider or platform is alleged or implied. Figures attributed to individual institutions derive from Group-IB detection telemetry and have not been validated by those institutions.

05

Readers are permitted to download, review, and cite this report for personal or educational use, provided that proper attribution is given to Group-IB and a link to the original source is included. Any commercial use or redistribution of the report without prior written consent from Group-IB is strictly prohibited.

06

This report is protected under applicable copyright and intellectual property laws. All trademarks referenced herein are the property of their respective owners and are used solely for informational purposes, without any implication of affiliation or sponsorship.

07	Unauthorized reproduction, distribution, or other use of the report, in whole or in part, without written consent from Group-IB, may result in legal action, including claims for damages.
08	Nothing in this report constitutes legal advice. Statements concerning the possible legal position of any party, including persons who paid for discounted settlement of obligations, are general in nature; readers should obtain advice from a qualified adviser in the relevant jurisdiction.
09	Indicators of compromise published in this report reflect observations at the time of research and may subsequently have been reassigned, recycled or remediated. Readers should not actively interact with the listed infrastructure and should independently validate any indicator before operationalizing it.
10	Any person or entity referenced in this report who considers that the reference is inaccurate or otherwise warrants review is invited to contact Group-IB through the contact channel on our website. The publisher will consider good-faith correction requests on their merits.

02 Acknowledgements

We would like to express our sincere gratitude to all teams whose expertise and dedication made this report possible. This investigation was led by Group-IB's Fraud Protection team, with contributions from specialists in fraud intelligence, threat research, and cyber investigations. Their collaborative work in analyzing and correlating phishing infrastructure, account-takeover activity, transaction telemetry, and cash-out marketplace intelligence was instrumental in uncovering the different layers of this scheme.

We also extend our appreciation to the financial institutions that supported the investigation and contributed to the identification of fraudulent activity. Their cooperation highlights the importance of intelligence sharing and cross-functional collaboration in addressing complex, multi-stage fraud.

Authors



Yuan Huang
Global Fraud Intelligence Lead



Andrei Loshchev
Team Lead, Fraud Consulting Analyst (Research), META



Rand Agha
Cyber Fraud Analyst, META



Tabish Amjad
Cyber Fraud Analyst, META



Abzal Aitoriyev
Junior Cyber Intelligence Analyst, META



Mahmoud Gamal Rashied
CERT Tier 2 Analyst, META

03 Introduction

Over the past several years, the Gulf Cooperation Council (GCC) region has invested heavily in payment security and fraud prevention. Mandatory 3D Secure on card payments, national digital-identity and authentication systems, real-time transaction monitoring, and close cooperation between banks and fintech providers have made the ecosystem markedly harder to abuse. Those investments have produced real results: many of the simpler methods fraudsters once used to cash out stolen cards no longer work as before.

One measure of that success is how much harder fraudsters now have to work. As the easy routes have closed, fraudsters have increasingly turned to more complex, multi-stage schemes that span multiple domains and exploit gaps between otherwise effective security controls. Group-IB's Fraud Protection team, working with banks in the region, identified one such scheme, in which stolen card details are used to settle genuine government bills and fines through official payment portals protected by 3D Secure (3DS). Approximately 300 such incidents were detected and flagged across several major retail banks.

The investigation illustrates the importance of Cyber Fraud Fusion (CFF): bringing together cyber threat intelligence, fraud intelligence, digital risk signals and investigation to build visibility across the full fraud chain. In this campaign, the individual stages span multiple environments that are typically monitored separately. Connecting those signals provides a more complete view of how the scheme operates—from initial victim acquisition and credential theft through account compromise, fraudulent payment and eventual monetization.

What makes the scheme notable is not any single clever step, but the way it is assembled across domains that rarely share signals. A paid search ad delivers the victim, a phishing kit harvests the card information, a mobile-operator eSIM swap can compromise the phone number, a bank account is taken over, a government portal receives an authenticated payment, a Telegram marketplace recruits the person who benefits, and cryptocurrency can provide an exit rail. Each domain sees only a fragment, and each fragment may appear unremarkable in isolation.

Rather than necessarily breaking individual payment controls, the scheme shows how fraudsters can combine compromised identities, devices and accounts to satisfy or work around them. In the cases examined by Group-IB, fraudulent transactions were completed with valid 3DS authentication, demonstrating why transaction-level controls alone may not reveal the preceding compromise.

The economics reinforce the disguise. Earlier attempts to cash out stolen cards in the region relied on topping up e-wallets before withdrawing the funds, but the tight integration between wallets and banks, together with the region's monitoring, made that route more exposed to detection and tracing. From October 2025, the operators shifted to settling real government bills instead. This produces what is best described as a bill-discount cash-out. The operator offers to settle a customer's genuine bill at a large discount, collects the discounted amount, and pays the full charge on the official portal using a stolen card. The customer's bill is settled, while the fraudulent charge is made against the compromised cardholder.

04 Key Findings

01

Fraud that passes every control instead of breaking it.

In the confirmed cases examined, fraudulent transactions were completed with valid 3D Secure (3DS) authentication after fraudsters compromised the cardholder’s mobile number and banking account. The scheme demonstrates how attackers can exploit compromised authentication factors without defeating the 3DS mechanism itself.

02

Approximately USD 2.01 million in confirmed fraudulent payments.

Group-IB identified confirmed losses across 80 compromised cards linked to three government institutions, within a wider set of approximately 300 related incidents flagged between October 2025 and August 2026. The scheme illustrates the significant financial impact that can result from the abuse of legitimate payment infrastructure.This is not a fringe technique; it is a large and growing share of card fraud.

03

Stolen card value is converted into discounted government bill payments.

Fraudsters offer to settle genuine bills or fines for members of the public at a discount, pay the full amount using a stolen card, and collect the discounted amount from the customer. The method exploits the trusted status of official payment portals while leaving the cardholder with an unauthorized charge.

04

Your existing controls cannot see it alone.

The fraud is only visible when a valid payment is correlated with the account-takeover signals on another channel. Single-channel monitoring, however good, allows it through by design. Group-IB Fraud Protection correlated valid 3DS transactions with account-takeover indicators observed across mobile and web banking channels, including newly registered devices, GPS spoofing and eSIM-related anomalies. This visibility helped identify fraudulent activity that may be difficult to detect when payment and account-access signals are assessed separately.

05

The scheme relies on participants from the public seeking discounted settlements.

Participants may be unaware of the illicit source of the funds or may knowingly participate in the arrangement. In either case, the payment is linked to their genuine government obligation, creating potential financial and legal exposure depending on the circumstances.

06

The scheme operates across three functional layers.

Group-IB assesses a phishing layer that acquires victims, a checker layer that validates credentials and takes over accounts, and a cash-out marketplace that recruits customers for discounted settlements. The evidence supports links between these layers, with low to moderate confidence whether they are operated by a single organization or cooperating specialist groups.

05 Attribution

Group-IB Threat Intelligence & Fraud Protection Portal

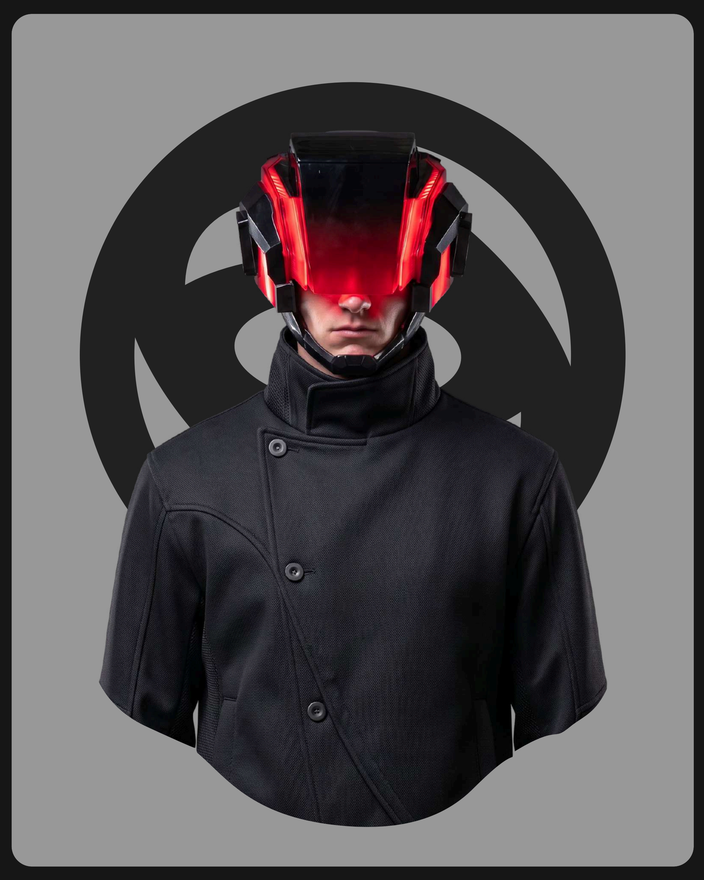
Group-IB customers can access our Threat Intelligence and Fraud Intelligence portals for more information on the actors, infrastructure and scam profile referenced in this research.

Threat actor profiles

The operation separates cleanly into three functional layers, each with its own evidence base: the phishing operator that acquires victims, the checker that validates and takes over accounts, and the cash-out layer that recruits the paying counterparty and settles bills.

Layer 1 Phishing Operator (Acquisition Layer)

The front end is run as a phishing service that impersonates mandatory government services. Group-IB CERT catalogued more than 400 resources across roughly ten disguise patterns and five distinct kits, promoted through verified Google Ads accounts. Group-IB’s Graph analysis links the clusters to common operators through shared identifiers. One kit domain is registered in Jordan, which is the only current infrastructure bridge to the checker layer's Ramtha, Jordan geolocation.



PAID MIRAGE

Language
N/A

First seen
Phishing activity observed through 2025

About

Operates a toolkit of phishing kits impersonating mandatory government service and online insurance aggregators, driven by geo-gated paid search ads. Harvests identity, card data, card PIN, OTP and the approval needed to issue an eSIM.

Aliases

- Phishing Operator (Acquisition Layer)

Skillset

- Phishing-kit development (AI site-builders, hand-coded and Bootstrap/PHP)
- Paid-search operations
- Real-time social engineering

Toolset

- Multiple kit families
- Telegram-bot, cloud-worker, Realtime-WebSocket and Pusher exfiltration
- CAPTCHA, geo-fencing and bot-scoring evasion.

Targeted Industries

- Mandatory government service
- Insurance aggregators
- Bank card brands

Motivation

- Financial gain

Modus Operandi

Places sponsored ads above genuine services for time-sensitive tasks; runs cloned booking or quote flows; several kits are live operator panels that steer victims in real time. Geo-gates content to the GCC area and serves decoys to researchers.

Layer 2

JORDAN CHECKER Group (Account-Takeover Layer)

The checker validates that stolen credentials work and removes the constraints on them. Group-IB Fraud Protection telemetry showed consistent device, network and geolocation characteristics across the sessions used to access victims' accounts, including eSIM use, GPS spoofing and low-cost iOS and Android devices. Correlating this evidence pointed repeatedly to a cluster of geohashes in Ramtha, Jordan, which is the basis for the group's name. A phishing domain registered in Jordan provides a second, independent link to the same geography.



JORDAN CHECKER

Language
N/A

First seen
October 2025 (government-portal abuse); peak January 2026

About

JORDAN CHECKER is the middle layer of the scheme. It validates that stolen credentials work, takes over the victim's online banking account, and removes the limits that would otherwise cap a fraudulent payment.

Aliases

- Checker Group in Ramtha Jordan region

Skillset

- Device-parameter manipulation and fingerprint spoofing
- GPS spoofing and eSIM roaming abuse
- Account recovery and OTP interception
- Mobile banking application operation at scale

Toolset

- GPS-spoofing tools
- Newly provisioned low-cost iOS and Android devices (iPhone 13 and iPhone 11 most common)
- Fraudulently issued eSIMs
- Mobile banking applications

Targeted Industries

- Retail banking
- Government payment portals
- Mobile network operators

Motivation

- Financial gain

Modus Operandi

The operators combine phishing-sourced card data and card PIN with a fraudulently issued eSIM for the victim's number. Because an eSIM receives calls and SMS exactly as a physical SIM does, the operator receives the victim's one-time codes. The observed takeover installs the bank app on a fresh device, masks location with GPS spoofing and eSIM roaming, starts account recovery using the stolen card data and PIN, intercepts the OTP through the hijacked number, and sets a new password. In 90% of investigated cases telemetry recorded a login from a previously unseen device fingerprint shortly before the fraudulent 3DS attempt. iOS dominated at 83% of sessions across 209 unique fingerprints, against 17% across 42 Android fingerprints.

Layer 3

CIVIC DRAIN (Cash-Out Layer Operators)

The cash-out layer recruits the paying customer and converts stolen card value into clean funds through government-portal settlement. Group-IB Threat Intelligence team found two linked Telegram channels, a recruitment storefront and a 2D/3D processing channel, that resolve to a single operator contact. Display names in both Arabic and Chinese were observed across the operator accounts.



CIVIC DRAIN

Language
N/A

First seen
Telegram storefront and processing activity observed from October 2025

About

CIVIC DRAIN is the cash-out layer. It recruits a paying member of the public and converts stolen card value into clean funds by settling that person’s genuine government bill through an official payment portal.

Aliases

- Cash-Out Layer Operators
- Bill-Discount Settlement Network
- “loaders” and “brokers”

Skillset

- Underground-market operation and customer recruitment
- Social engineering
- Carding logistics and payment-gateway procurement
- Cryptocurrency settlement

Toolset

- Paired Telegram recruitment and processing channels
- 2D/3D payment links on regional gateways
- USDT-TRC20 wallets
- Escrow and referral arrangements

Targeted Industries

- Government fines, utility and property billers
- Retail banking
- The paying public as counterparties

Motivation

- Financial gain

Modus Operandi

The market runs as two cooperating Telegram channels under a single operator contact: a public recruitment storefront advertising discounted settlement of utility bills and traffic violations, and a processing channel selling "2D and 3D links" on named payment gateways. Pricing is a fixed, openly advertised rate card rather than a negotiation, ranging from roughly 20% to 50% of the bill. Operators require payment upfront in USDT-TRC20, describe their funding source as "cards with balance that need draining", and openly describe the business as money laundering. A consistent role vocabulary - loader, broker, receiver, customer - indicates a specialised division of labour rather than isolated actors.

One group, or several?

The three layers map one-to-one onto the three halves of the attack chain, and the evidence supports each layer on its own. The open question is whether they are a single organisation running the whole operation or up to three specialist groups cooperating. The coupling between the layers is not uniform, so the layers are best kept distinct while the linkages are stated honestly.

The phishing and checker layers are more likely linked. Two bridges connect them: a phishing kit domain registered in Jordan matches the checker's geolocation in Ramtha, Jordan, and the eSIM-swap is set up inside the phishing flow (the victim approves the eSIM issuance on the phishing page) and then exploited by the checker. That is a direct handoff rather than proof of one operator, but it is stronger than the convergent method alone.

The link between the checker and the cash-out layer rests on convergent evidence rather than a shared identifier. The two operate against the same scheme, the same targeted institutions, the same victim geography, the same timeframe and the same value profile. They are also operationally complementary: the checker produces compromised cards and authenticated portal access, and the cash-out market is built to monetise exactly that kind of access.

The available evidence supports links between the phishing, checker and cash-out layers, but does not establish that they are operated by a single organization. Group-IB assesses that the phishing and checker activity is more closely connected, while the relationship with the identified cash-out marketplace is supported by convergent operational indicators. The possibility of a single organization or cooperating specialist groups remains open, and the available evidence does not conclusively attribute the confirmed losses to specific Telegram operators.

Scam profile

Government-Portal Bill-Discount Cash-Out

GROUP-IB

Attackers leverage SMS phishing campaigns impersonating trusted courier and postal services. Victims receive urgent delivery notifications containing malicious tracking links that redirect to counterfeit shipment websites designed to harvest personal and financial information. By using disposable domains, sender ID spoofing, and mobile optimized phishing pages, attackers can scale this campaign globally while maintaining credibility and reducing detection.

First seen:

7 Mar 2025

Techniques:

- Darcula phishing as a service platform
- Phishing kit mechanics
- SMS delivery mechanisms

Targeted Industries:

- Postal Services
- Financial Services

Targeted Countries:

- GCC

Victim Path

01 Victim receives phishing SMSes containing phishing links.

02 Links lead to convincing fake tracking pages, optimized for mobile browsers.

03 Victim believes they are resolving real delivery issues, fills in personal details.

04 Embedded scripts begin real-time keylogging.

05 Data is transmitted to attacker-controller servers via persistent WebSocket connection.

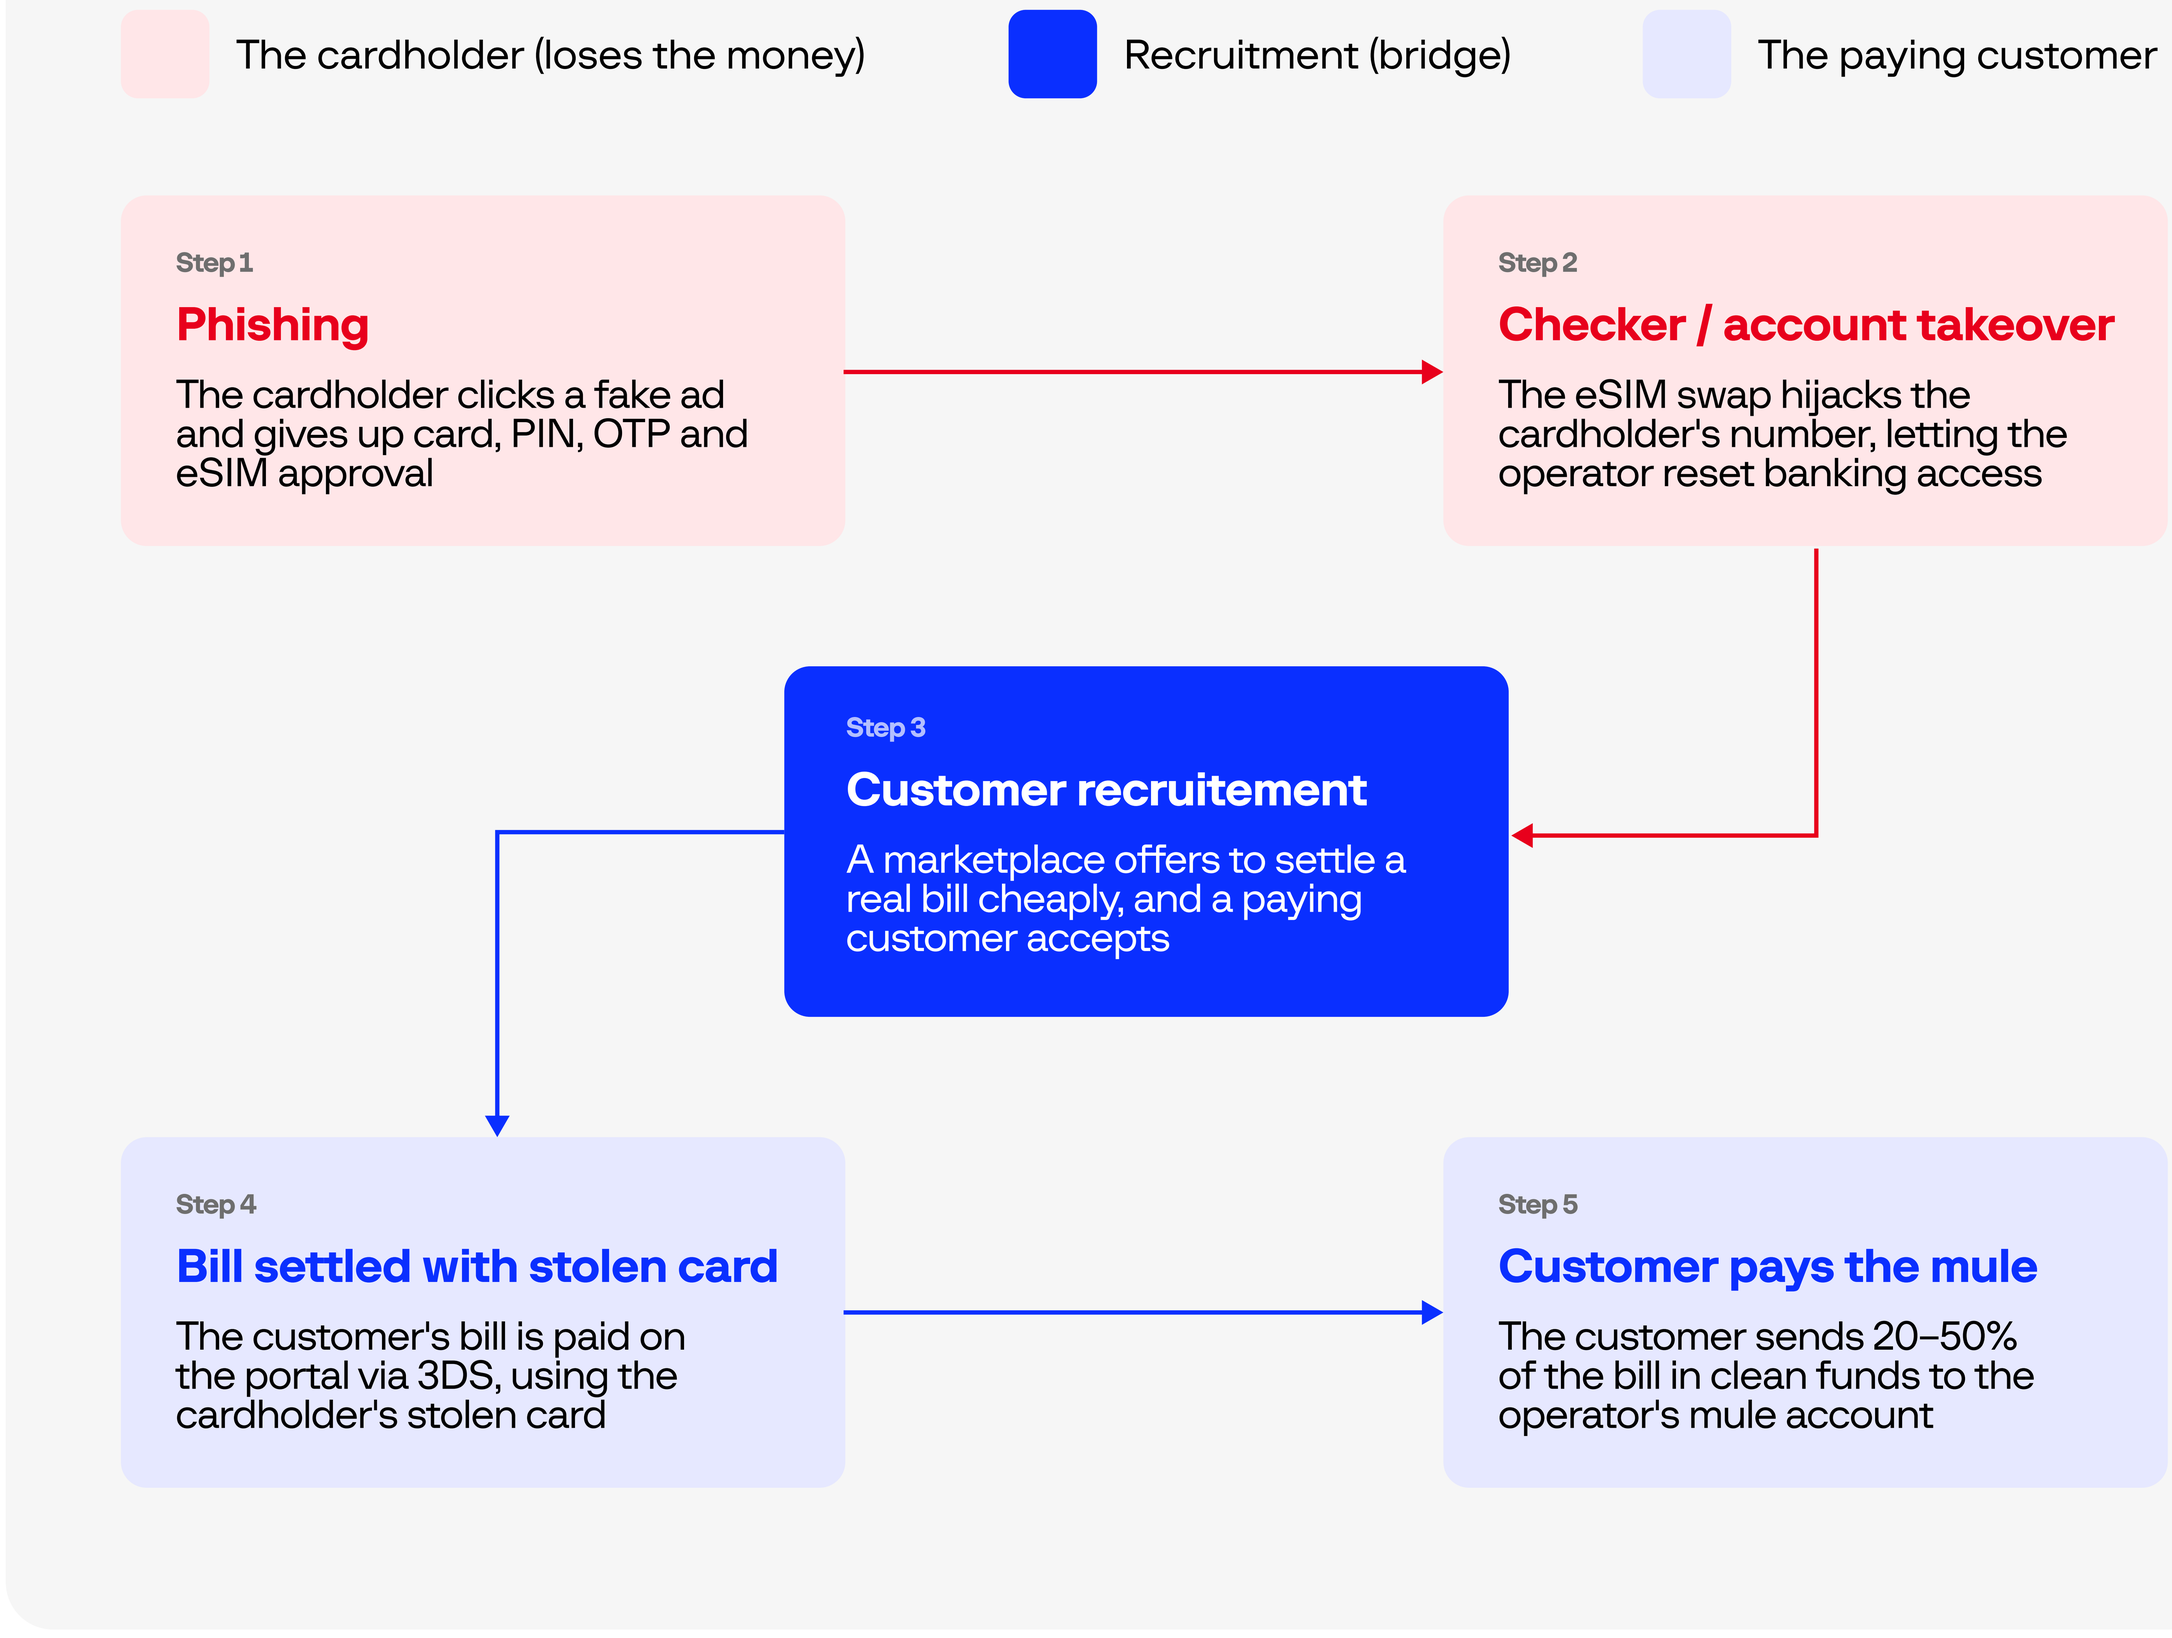
06 Fake payment page then requests a token re-delivery fee.

07 Victim enters banking and credit card info.

08 Funds are stolen and payment credentials exfiltrated via the same WebSocket channel.

Group-IB, 2026

Victimology



The Scheme, step by step

The confirmed financial impact concentrated across three GCC government institutions. The institutions below are payment channels abused by the operators; no failing on their part is alleged or implied. Figures are derived from Group-IB detection telemetry for 80 compromised cards during the observed period and have not been validated by the institutions themselves.

Institution	Role in the scheme	Confirmed fraudulent value
Justice-sector institution	Settlement of legal fines (69% of fraudulent transactions)	USD 1,385,100
Electricity institution	Settlement of utility bills	USD 561,323
Property and rental institution	Settlement of property charges	USD 66,909
Total	Three billers, across 80 compromised cards	Approximately USD 2,0 million

Around 300 incidents were flagged in total. The confirmed figures above are the subset validated to these three institutions.

06 Government Payment Cash-Out Scheme

From wallet top-ups to government portals

A shift in cash-out patterns is visible in the banking data. Through most of 2025, fraudulent 3DS sessions clustered on non-government merchants and fintech top-ups. From October 2025, the average fraudulent transaction value on government portals rose sharply, indicating the emergence of the government-payment cash-out scheme observed in this investigation. Activity in the analyzed dataset peaked in January 2026.

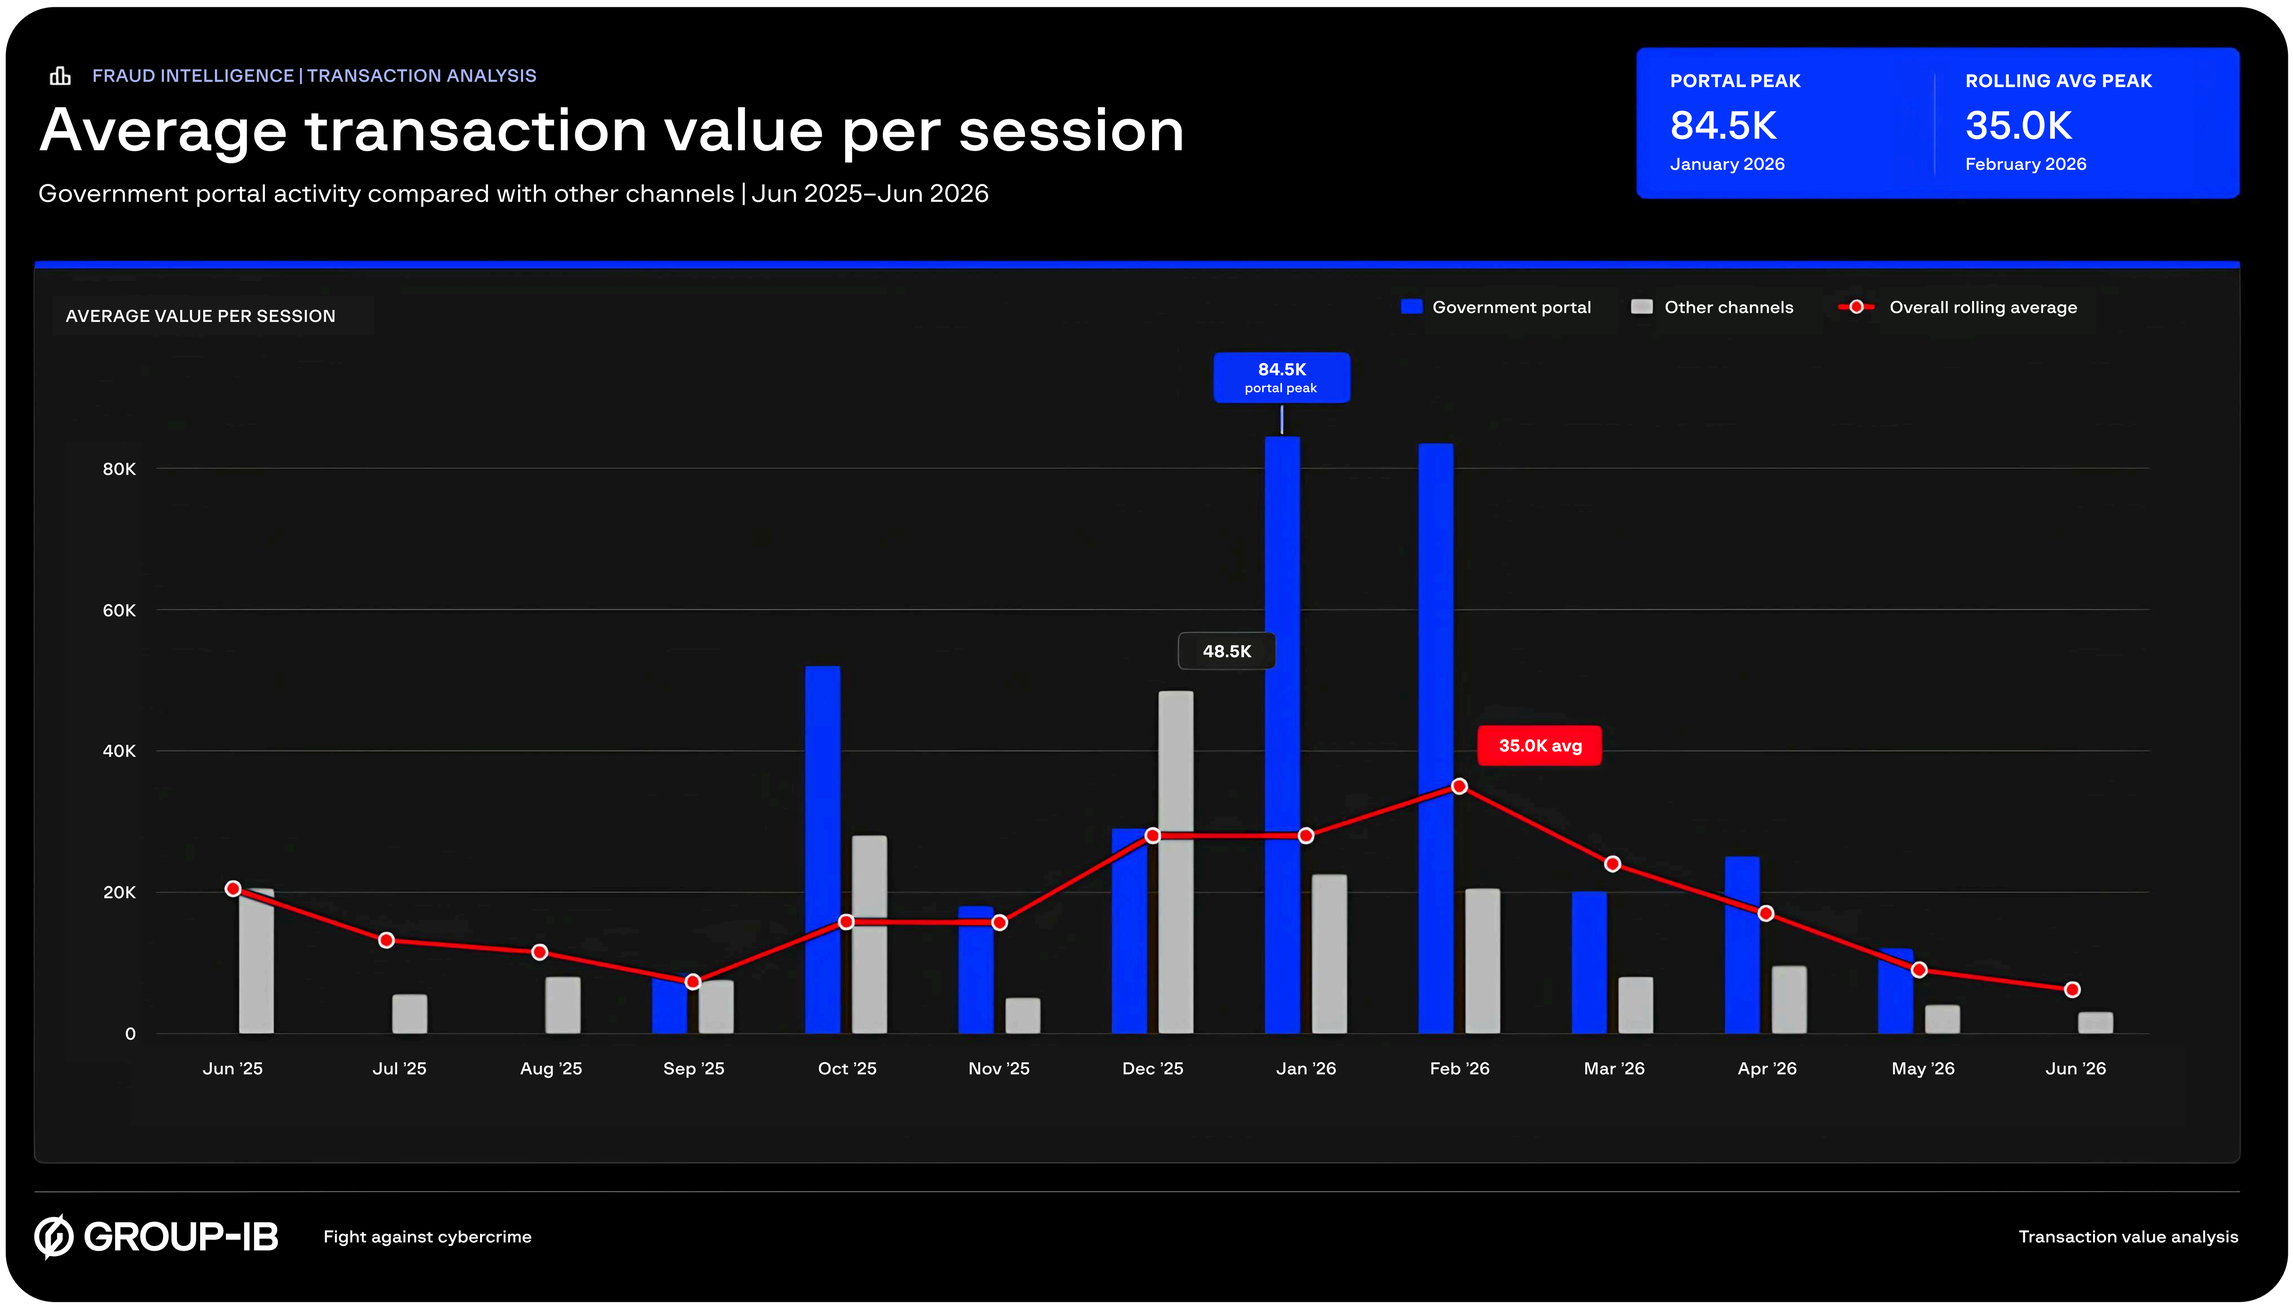


Figure 1. Average transaction value per fraudulent 3DS session, non-government vs government portals over time.

The logic of the shift is straightforward. A payment to a justice-sector fines portal or to the electricity utility is the least suspicious thing an account can do. It is one-way, hard to reverse, and rarely flagged by monitoring tuned to detect payments going to unfamiliar or high-risk recipients. By routing stolen value through a trusted government biller on behalf of a real person, the operator ensures that no account belonging to the fraudster ever appears in the transaction, so there is nothing obviously suspicious to trace.

Two victims, defined

Because the cash-out is the transaction that links them, the two parties must be distinguished clearly:

Victim A, the cardholder.
The person whose card data, identity and online-banking account are compromised through phishing and takeover. Victim A loses the funds.

Customer, the paying counterparty.
A member of the public who buys a discounted settlement of their own real government bill. The evidence indicates the customer may be aware the service is illicit: the operators openly describe the business as money laundering and state that “all the loaders work this way” with our analysts. The customer is therefore treated as a potentially knowing participant and remains legally culpable.

The attack chain

The operation runs as a staged chain in which each stage feeds the next. The front end acquires victims and harvests credentials, the middle validates access and takes over the banking account, and the back end converts stolen value into clean funds through a trusted government biller.

Stage 1:

Acquisition and targeting

Victims are acquired predominantly through paid Google search advertising, geo-targeted to users inside the GCC. Group-IB CERT-GIB identified more than 70 keywords tied to the scheme; for every tested variation, a sponsored advertisement appeared at the top of the results, placing the fraudulent page above the genuine service. Victims click because they are actively searching for a mandatory, time-sensitive task, the ad uses the correct Arabic copy, and the display URL looks plausible. Two lure families were observed:

Mandatory government services.
The kits impersonate government service providers and clone the official booking portal or the responsible authority to add credibility.

Insurance renewal.
The kits impersonate online insurance aggregators, presenting a familiar “compare and buy” quote flow.

Both tasks are routine, recurring and time-pressured, which can lower users’ vigilance. Some ads lead straight to a single-page phishing site; others pass through a thin intermediary “booking” page whose two buttons forward the victim to the current live phishing destination. This design allows operators to replace the phishing destination when it is blocked, while keeping the ad and intermediary landing page active.

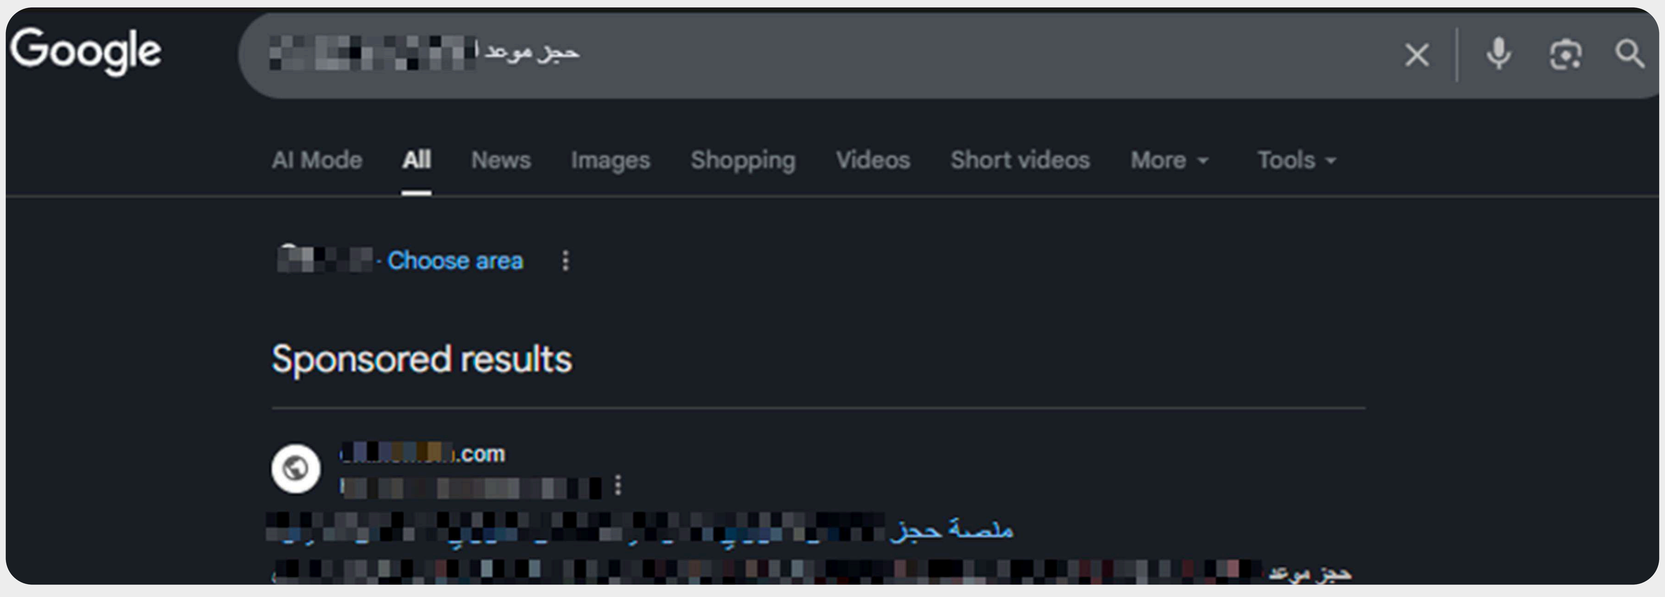


Figure 2. A sponsored (paid) Google ad for a fake government service booking site



Figure 3. The intermediary "booking" page shown between the ad and the phishing site

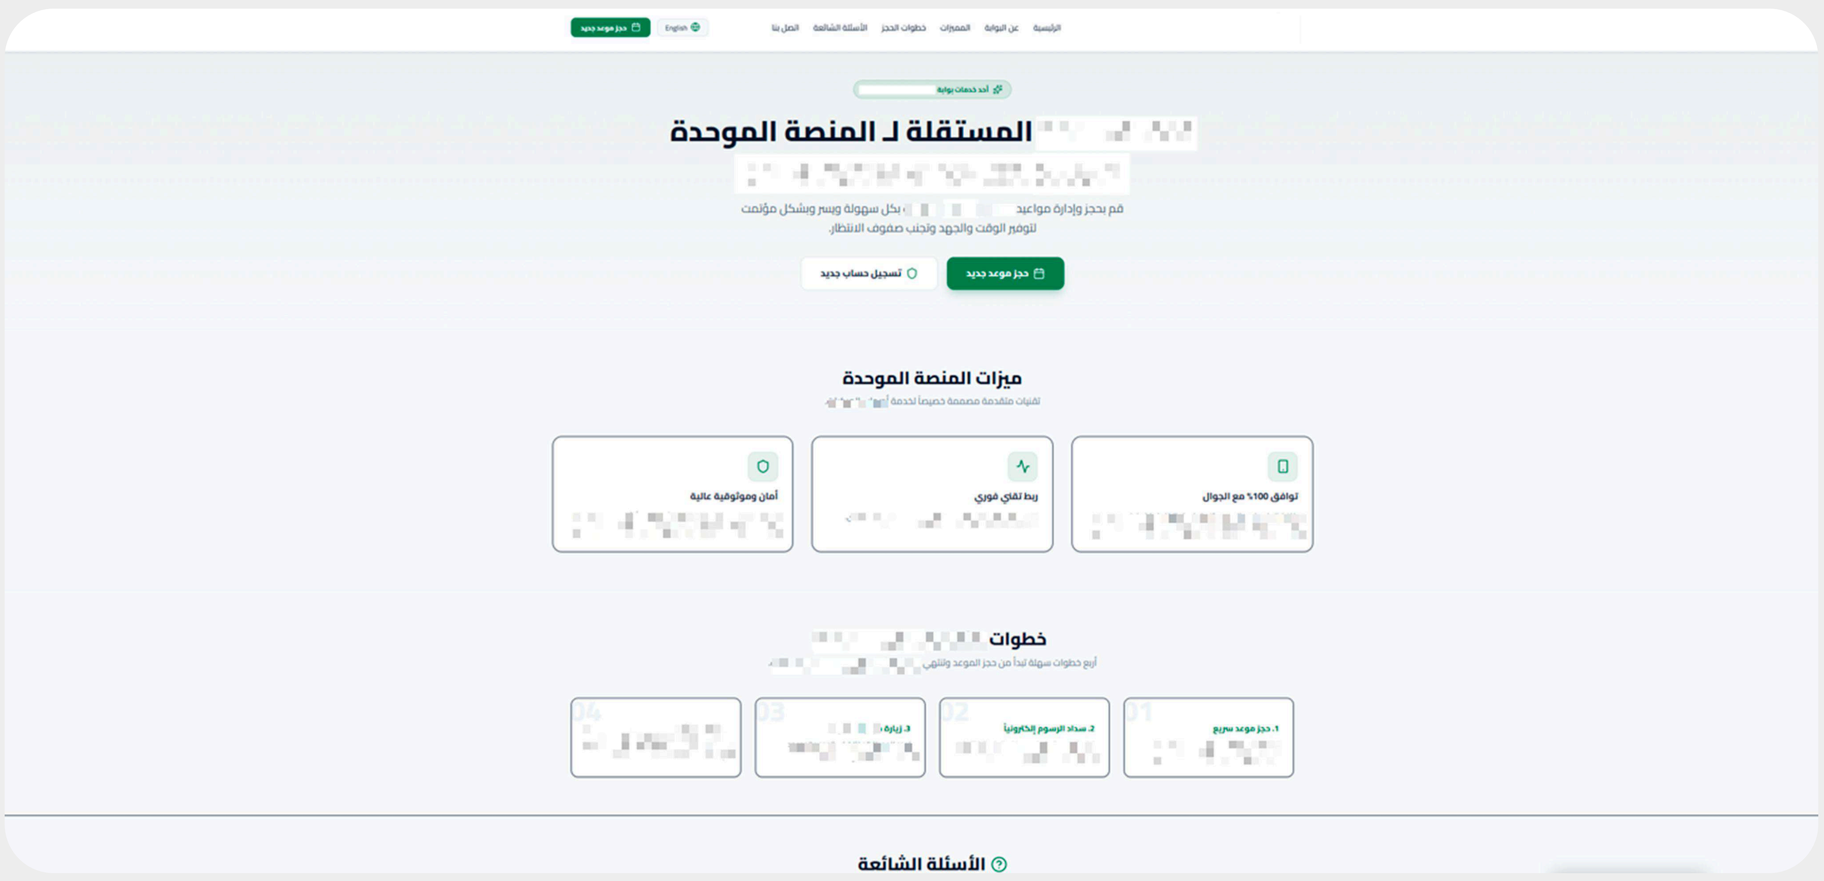


Figure 4. Fake Unified Platform for Government Service Appointments

Stage 2

The phishing toolkit

The front end is not a single fake website but a toolkit of technically distinct kits, more than 400 resources across roughly ten disguise patterns, built on different platforms and wired to different collection channels.This diversity helps spread operational risk, complicates signature-based blocking, and allows operators to rotate infrastructure. The kit families are summarised below, with the purpose each design serves. Full per-kit fingerprints are provided for defenders in the IOC section.

Kit family	Build and hosting	Collection channel	Why it is built this way
AI-generated SPA	React via an AI site-builder, disposable domains	Attacker-owned cloud workers	Fast to clone and rotate; the whole cluster ties to one builder account via a shared asset ID. Outside the targeted country, it just shows an endless loading screen.
“Map-enabled” SPA	React plus map widget behind a CDN	Live WebSocket streaming input as typed	Fake “nearest branch” map adds credibility; a reload trick defeats cached analysis
Authority impersonation	Vue/React on a hosted AI-app platform	Split to database plus Telegram, or Realtime C2	Falsely claims official supervision; runs the OTP no-exit loop
Static PHP-based kit	Bootstrap and PHP on free hosting	Server-side PHP endpoint	Cheap and disposable; behavioural bot-scoring blocks scanners
Hand-coded multi-page	Plain HTML/JS on free hosting	Telegram bot in cleartext	Simplest to stand up; collection channel visible in page source

Four design features matter most for defenders, because each has a clear purpose:

01 **Live operator panels.**
Several kits are not passive forms. As the victim types, their input is streamed to an operator who decides what to show next and can push a card page, an OTP page or a “please wait” screen on demand. This lets the operator react in real time to the authentication steps presented to the victim.

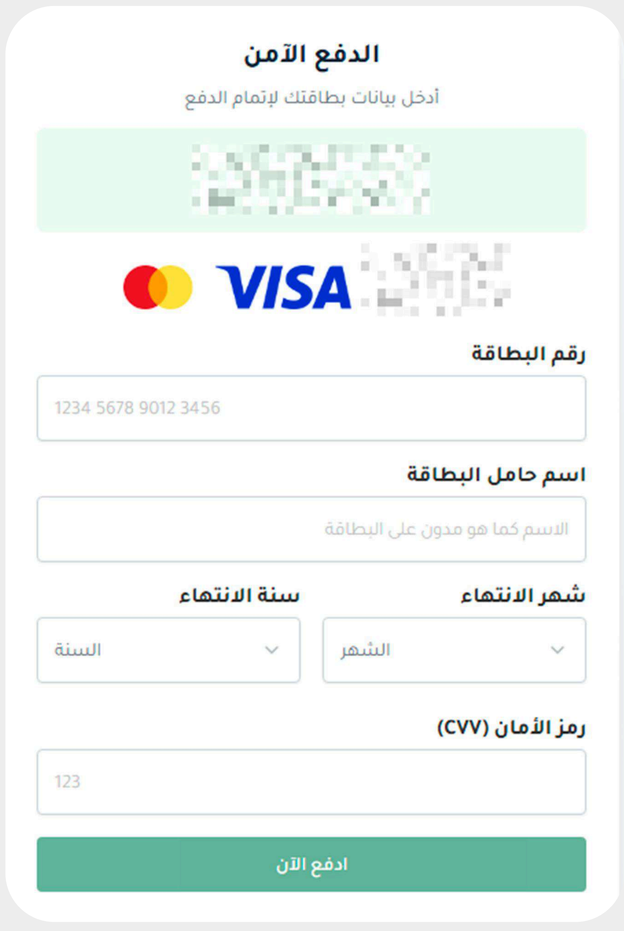


Figure 5. Fake “secure payment” page harvesting card number, cardholder name, expiry date and CVV.

02

Apple Pay disabled to force card entry.
The payment page displays both Apple Pay and card options to look legitimate, but selecting Apple Pay returns a “temporarily unavailable” error. This steers the victim onto card entry, which is the data the operator needs to harvest.

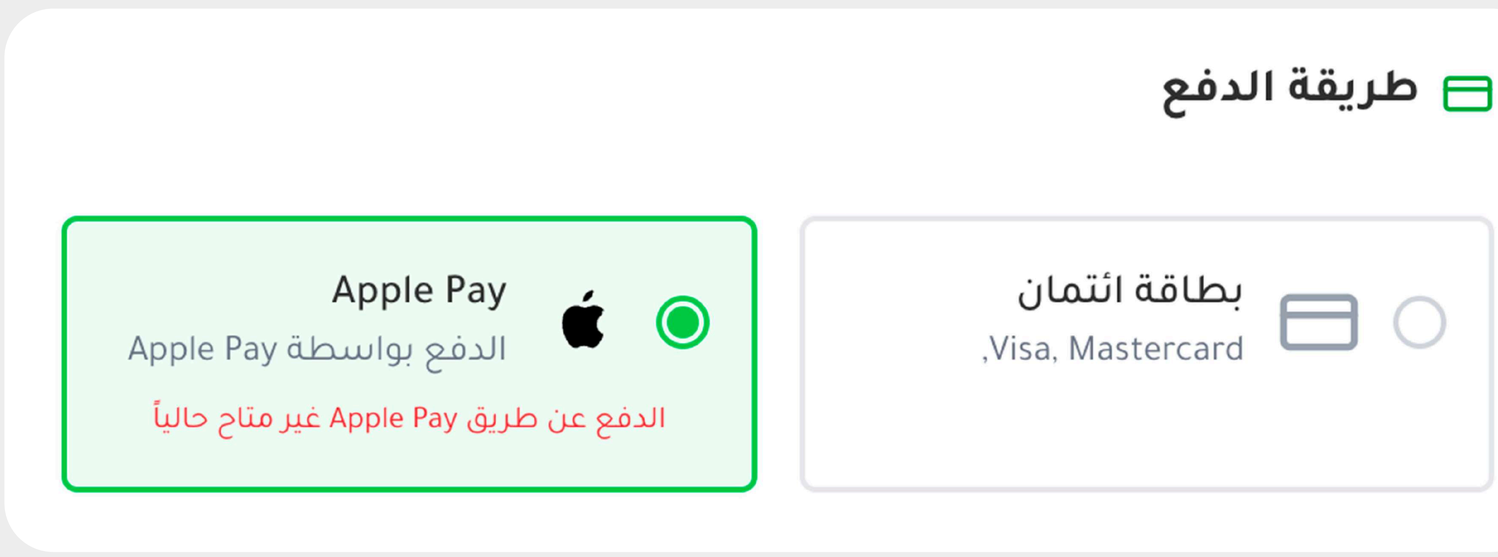


Figure 6. Apple pay displaying an unavailable message on the phishing domain.

03

Fake bank discounts and a countdown.
The card page shows fake cashback or discount promotions tied to named banks, with a live countdown, to pressure the victim and increase the chance they submit their card.



Figure 7. Fake “30% cashback” bank-card offer with a countdown timer.

04 An OTP no-exit loop.

In the examined kits, the OTP step never validates the code. It always reports the code incorrect and asks again, handing the operator a fresh, live authorisation code on every attempt. Operator code comments instruct the kit to forward the data “regardless of the code”.

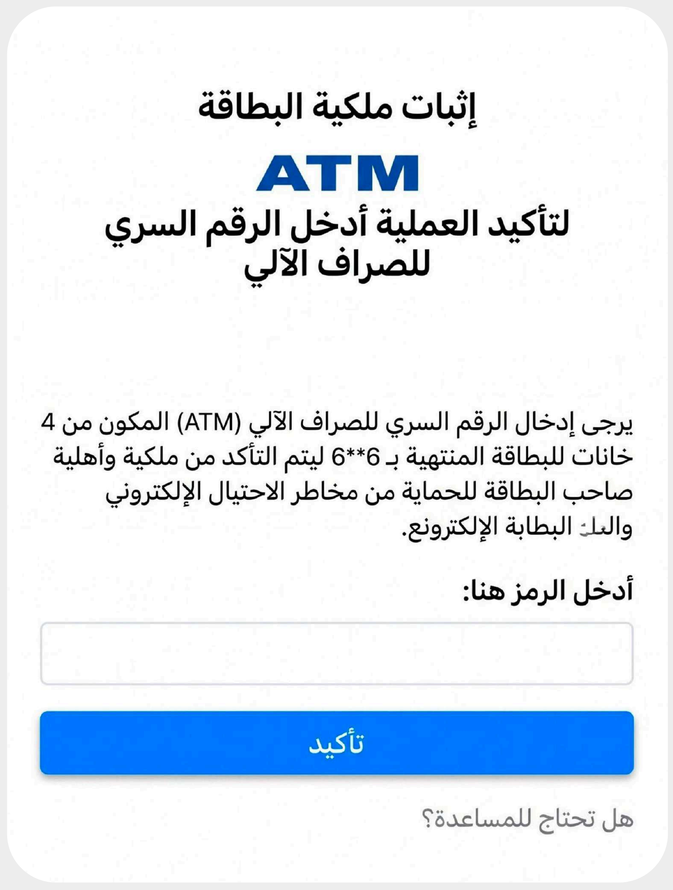


Figure 8. Fake OTP/ATM PIN page prompting the victim to enter a secret card PIN to “confirm the operation.”

Two further points are worth noting. The operator captures the victim's identity before the card step, so anyone who abandons payment has already surrendered enough personal data f that could be used to support a convincing follow-up call. And across the kits, layered evasion (geo-fencing to GCC, VPN and datacentre filtering, CAPTCHA, behavioural bot-scoring, and decoy “404” or “country not allowed” shells) helps keep the content accessible to targeted victims while limiting exposure to automated analysis and researchers.

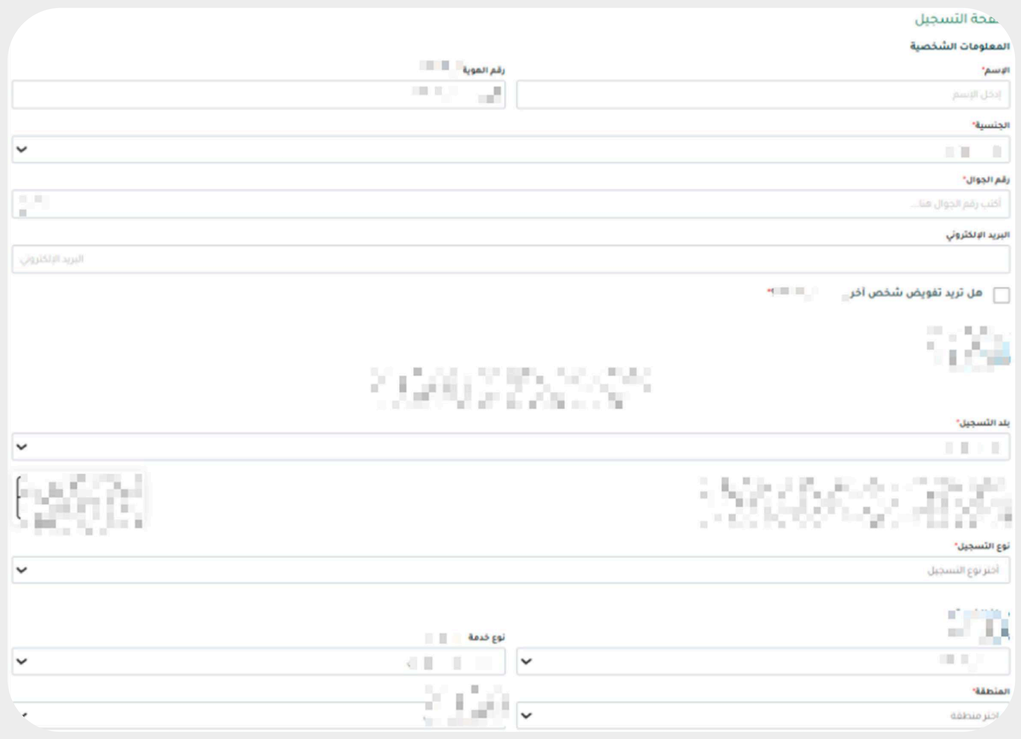


Figure 9. Fake registration form harvesting personal details related to the government service.

Stage 3

Credential and authentication harvest

Across the kits, three categories of data are collected: financial (card number, cardholder name, expiry, CVV and card PIN), personal (full name, national ID, registered address), and mobile-account data including the approval needed to issue an eSIM. A critical step in the observed attack chain is the verification flow. After “payment”, the victim is asked either to enter an OTP tied to their card account, or to approve a prompt in a national identity application. The prompt is presented as a request from a telecommunications regulator. In reality, the approval authorises an eSIM issuance request with the victim's mobile operator, enabling a SIM swap that can give the operator control of the victim's number for subsequent stages of the attack.

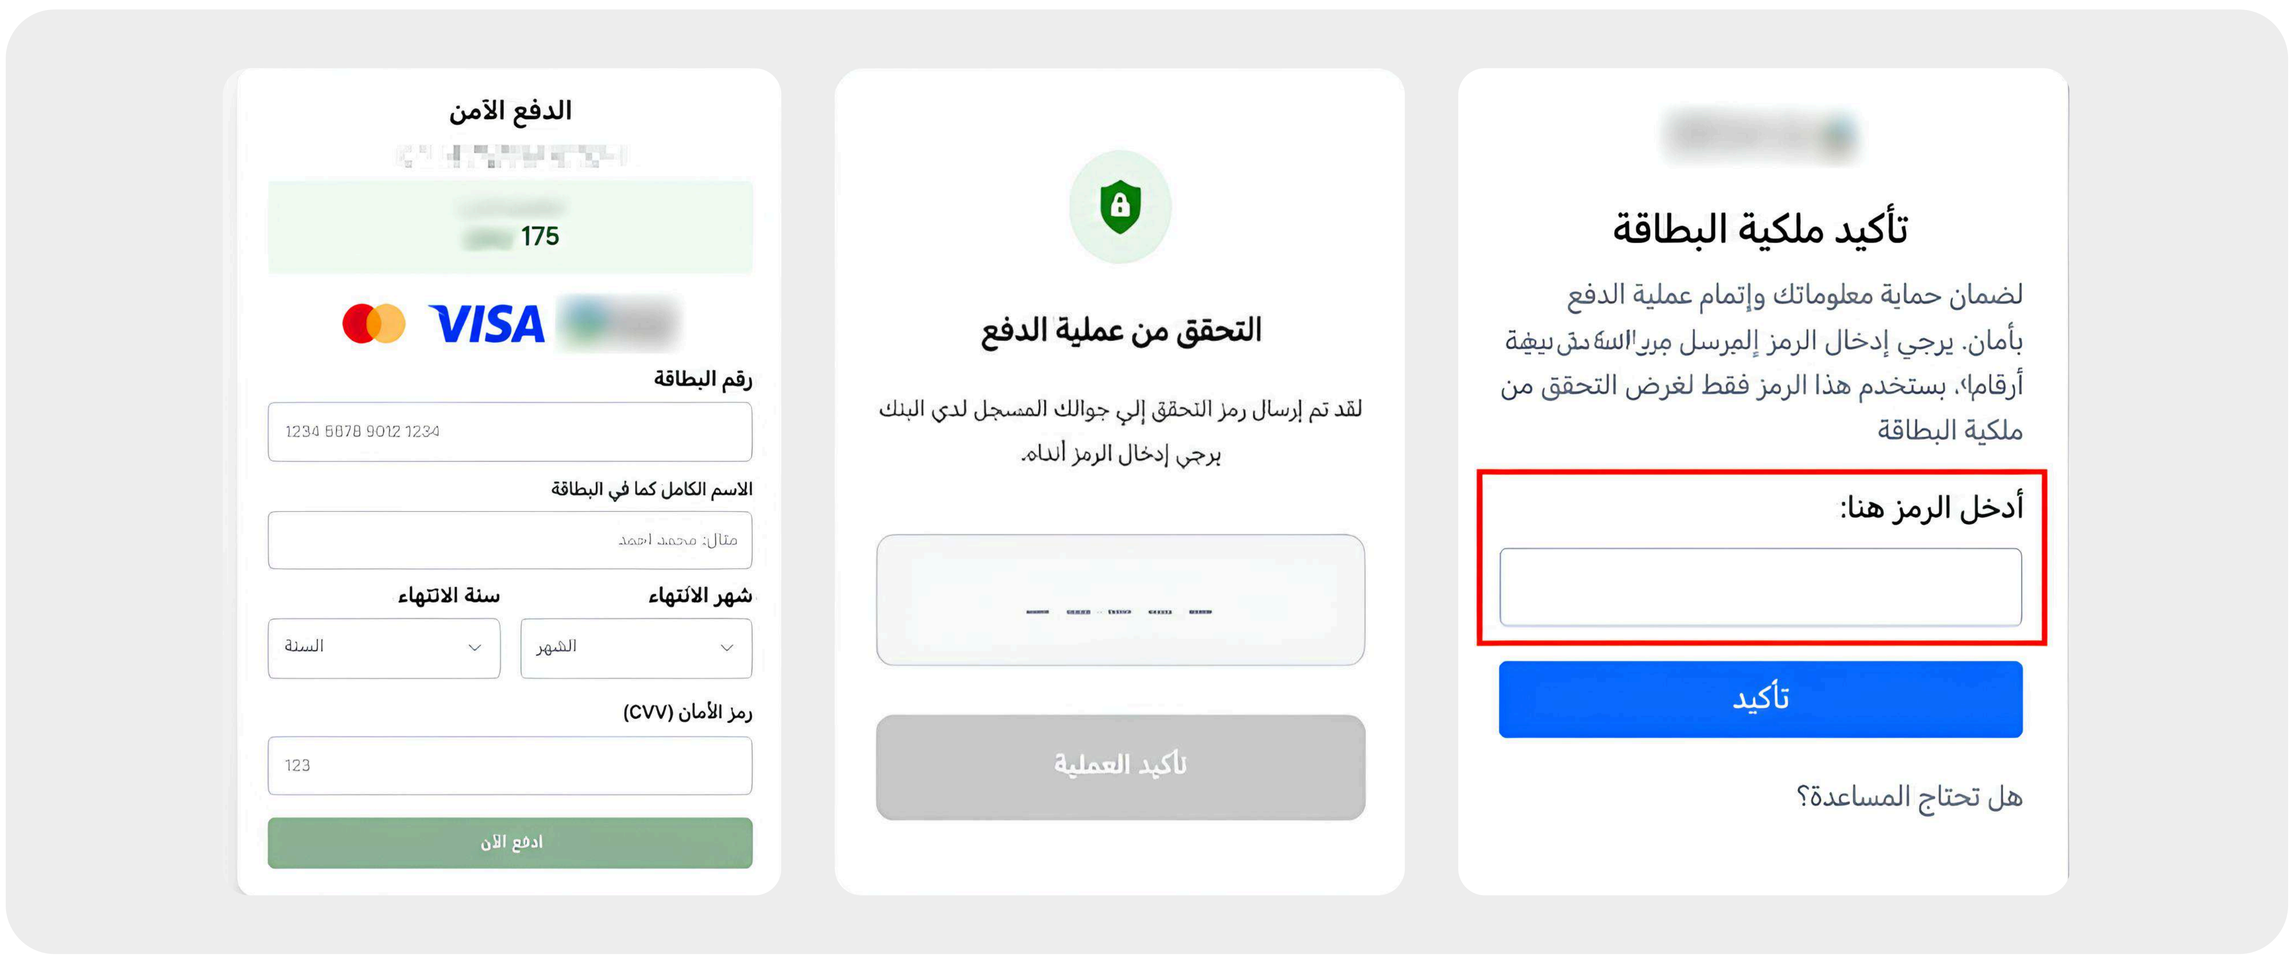


Figure 10. Phishing payment form requesting card number, cardholder name, expiry date, CVV, OTP and PIN.

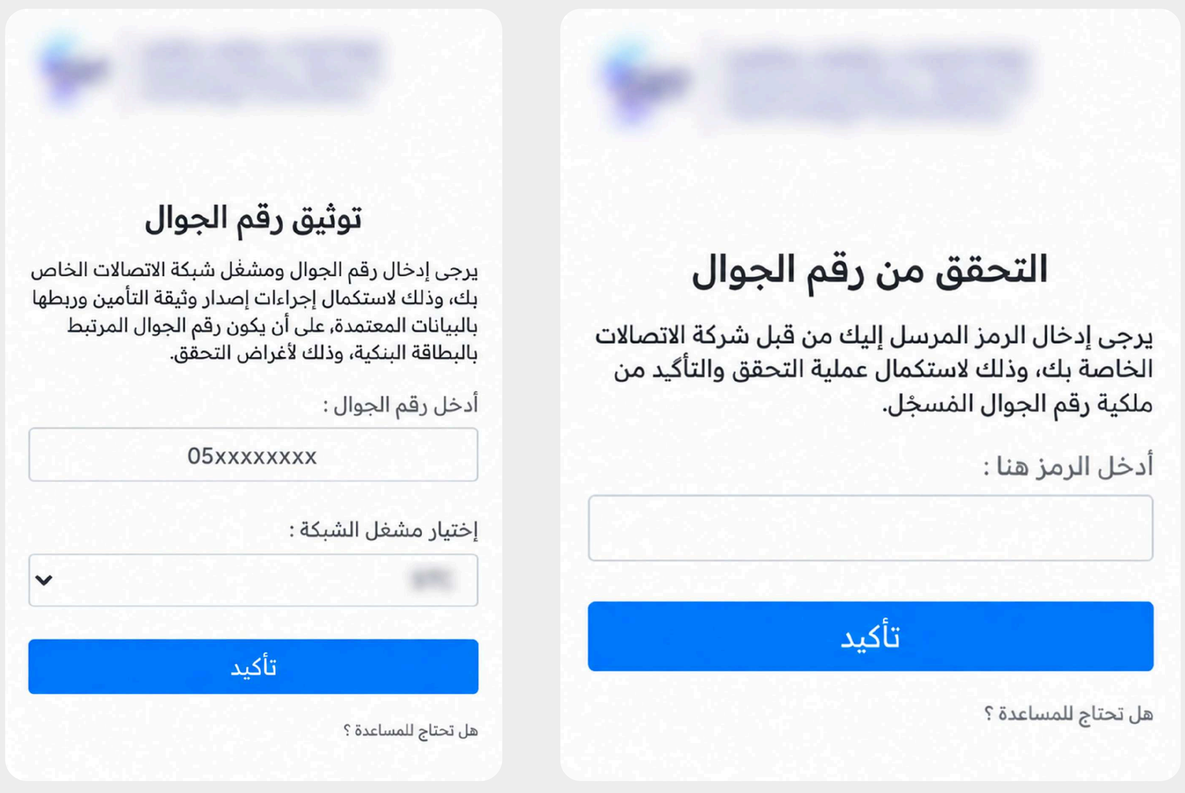


Figure 11. Phishing page imitating a national digital identity approval screen.

Stage 4

The checker layer (account takeover)

This is the stage the operators themselves call the “checker”, which validates that stolen credentials work and helps overcome the security controls protecting the account. Because an eSIM receives calls and SMS exactly as a physical SIM does, an operator who has fraudulently activated an eSIM for the victim's number can also receive the bank's verification codes. Where a bank permits online-banking recovery using card data, the card PIN and an SMS OTP, the whole chain of factors has already been compromised in a single campaign.

The observed takeover proceeds by installing the bank app on a fresh device, masking location with GPS spoofing and eSIM roaming, starting account recovery with the stolen card data and PIN, intercepting the OTP through the hijacked number, and setting a new password. Once inside, the operator may be able to raise card limits, move funds between the victim's own accounts, and approve 3DS challenges in-app, providing the capabilities needed to complete a high-value payment.



Figure 12. Group-IB Fraud Protection real time detection of Account Takeover (ATO) in mobile application.

Group-IB Fraud Protection detected that in 90% of investigated cases, telemetry recorded a login from a previously unseen device fingerprint, with indicators of GPS spoofing and eSIM-swap, shortly before the fraudulent 3DS attempt. iOS devices dominated: 83% of sessions across 209 unique iOS fingerprints, against 17% across 42 Android fingerprints. The iPhone 13 and iPhone 11 were the most common models. The shift between pre-attack and takeover locations frequently produced an impossible-travel pattern that provided investigative indicators consistent with activity originating from Ramtha, Jordan.

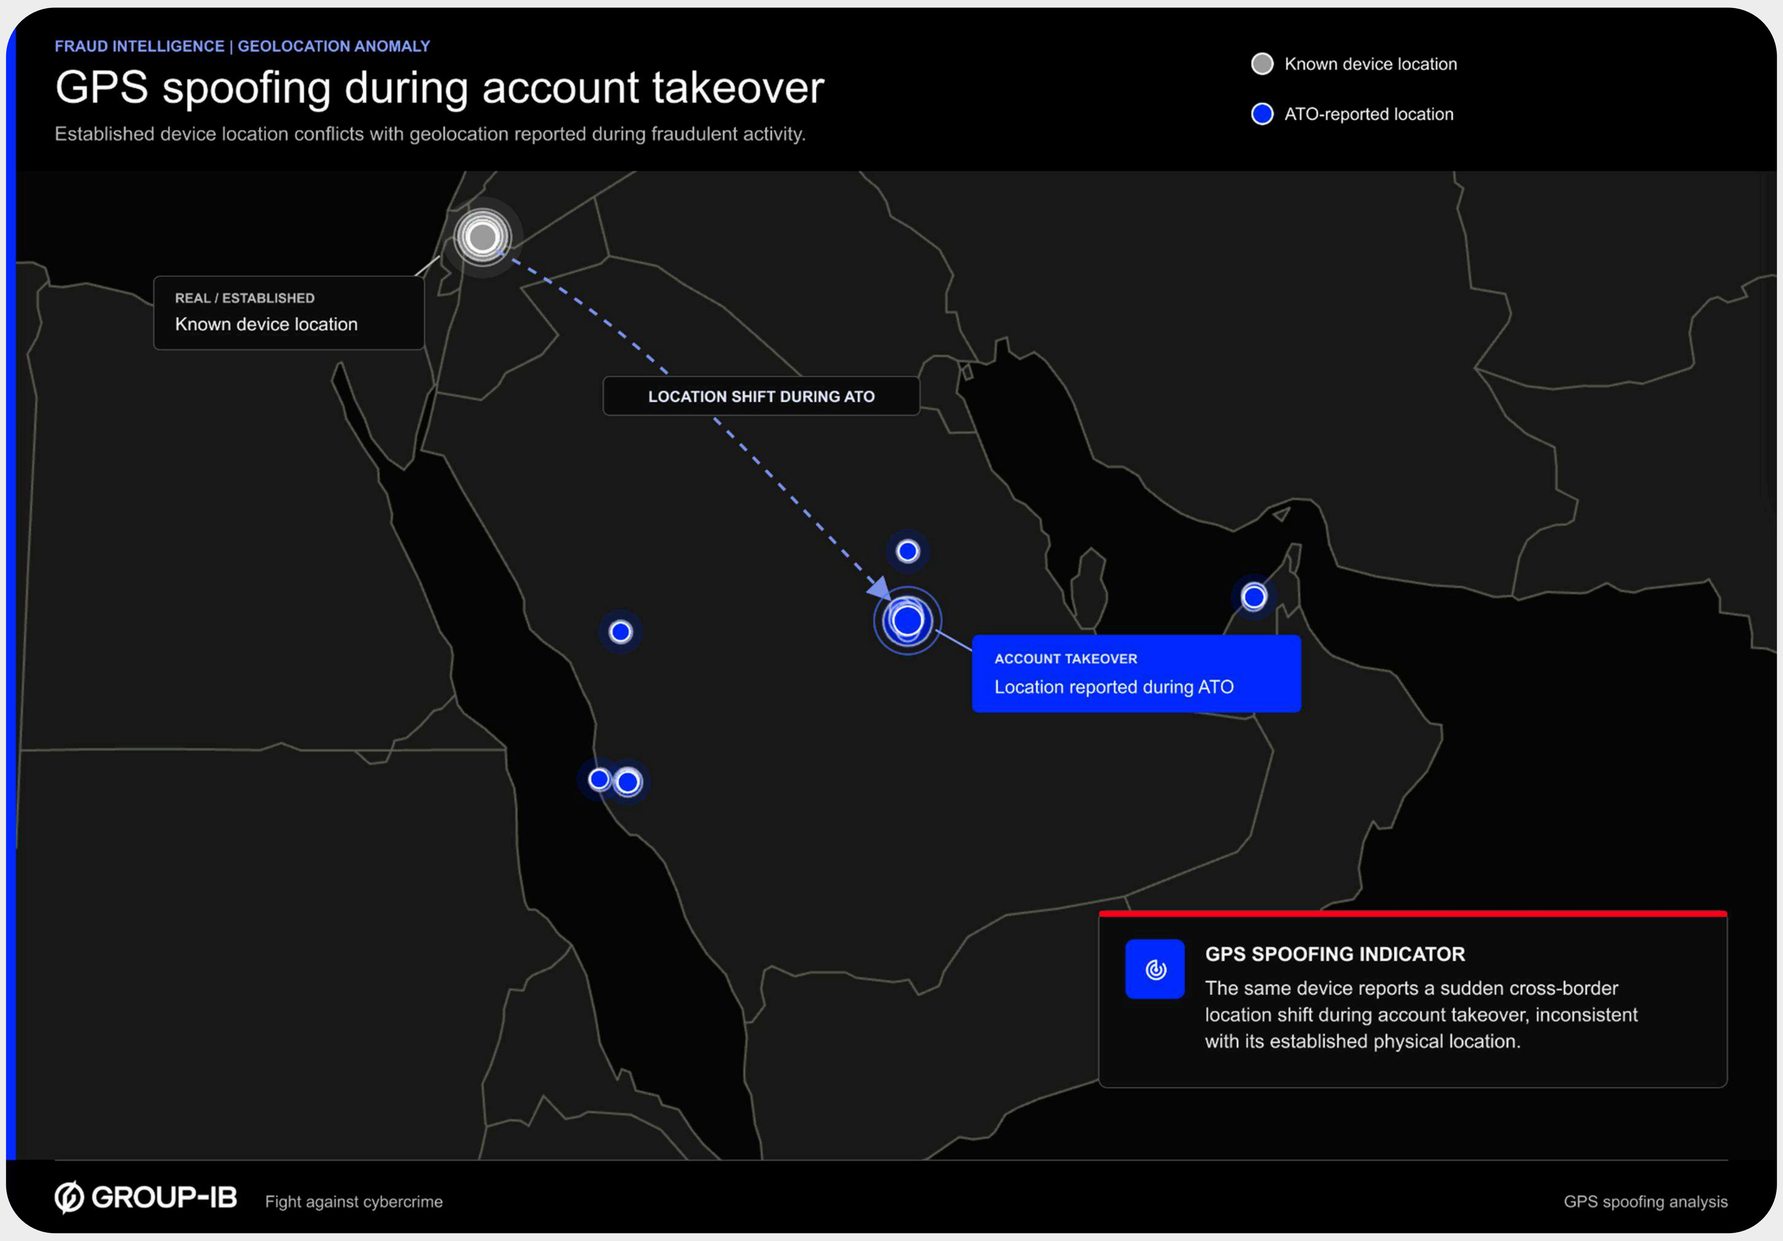


Figure 13. Impossible-travel and GPS-spoofing indicators observed during an ATO session.

Stage 5

The cash-out (government portal and customer)

After takeover, the operator no longer needs to defeat 3DS, because they can satisfy it. The purpose of 3DS is to authenticate the person authorising a payment, using methods that may include a one-time code or in-app approval. In the observed scheme, the compromise of the victim's phone number and banking access enables the operator to receive or approve authentication challenges. The issuer sees a correctly satisfied authentication. In our dataset, every single fraudulent transaction completed through a fully valid 3DS authentication. The control was performed correctly on a premise that had already been compromised.

The Central Mechanic

3DS is not broken here. It is answered correctly, by the wrong person, which can make the transaction appear less suspicious than a transfer to an unfamiliar recipient. The money then lands on a trusted government biller on behalf of a real person, which is far less likely to be flagged or reversed than a transfer to an unknown account.

The operator uses the authenticated session to settle the customer's real government obligation (a fine, an electricity bill) on the official portal using Victim A's card. The customer's debt is cleared, Victim A loses the full charge, and the operator keeps the spread. Because the funds move from Victim A to a legitimate government entity on behalf of the customer, the fraudster-controlled account does not appear as the recipient of the government payment.

Stage 6

The underground market and customer recruitment

The recruitment of the customer and the settlement itself are run as an organised Telegram market, and this market is one of the central findings of the research. It operates through two cooperating Telegram channels linked by a shared operator contact, uses advertised rate cards, and involves several specialised roles that together complete the cash-out.

Two cooperating channels: recruitment and processing

The person whose card data, identity and online-banking account are compromised through phishing and takeover. Victim A loses the funds.

Recruitment storefront.

A public-facing channel advertises discounted settlement of electricity bills, traffic violations and “all services for residents”. It names a single contact and carries the rule “deal only through the owner or the admins”. This is the storefront that finds the customer.



Figure 14. Example Telegram channel (391 members).

Processing channel.
A second channel (public-facing channel, around 155 subscribers) sells the charging capability: “2D and 3D links” on named payment gateways. A 3D link relays the 3DS challenge and captures the victim OTP in real time, which is the same OTP the eSIM swap targets. This is the layer that actually pushes the payment.



Figure 15. Telegram channel (155 members)

How we know the two channels are one operator

The linkage is supported by a shared contact. Both channels publish the identical operator contact handle (@A_L_i_H_N), which is the strongest single piece of evidence that recruitment and processing are coordinated through a common contact. Separately, one listed administrator of the recruitment storefront (display name “Darajoun”) was observed in a large Visa-card carding supergroup of around 1,400 members, under a Chinese-language display name meaning “Middle East cash withdrawal, VIP tools, escrow”, soliciting payment gateways and payment links “on condition the link owner is direct”. The link between the two accounts rests on a display-name match only, and Group-IB does not assert that they are operated by the same person. On the available evidence, recruitment, the card-processing layer and gateway procurement appear to converge on a single operator.

The roles that complete the cash-out

The market uses a consistent role vocabulary, which shows a specialised division of labour rather than isolated actors:

Role (Arabic term)	Function in the cash-out
Loader (لودر)	Settlement of legal fines (69% of fraudulent transactions)
Broker (وسيط)	Resells the loader's capability to the customer for a margin and can act as escrow to reduce the risk of either side being cheated.
Receiver (مستقبل)	Absorbs value into a mule account or wallet and forwards clean funds onward.
Customer (العميل)	Supplies clean money and receives a settled bill. The customer's legal exposure depends on their knowledge of the scheme and the circumstances of their participation.

What the customer pays: the discount model

Based on the observation, pricing is a fixed, openly advertised rate rather than a negotiation, and it was corroborated across an in-channel rate card and three direct-message engagements conducted under an authorised research persona (no funds were transacted and no real data was supplied). The customer pays only a fraction of the bill's face value, and the debt is cleared in full. The operator's economics are extreme: because the bill is paid with someone else's stolen card, almost the entire amount the customer pays is profit.

Source	Customer pays (share of the bill)	Bill types	Effective discount to customer
Operator “Dark Baron” (DM)	~20%	Fines and utilities (trust-ramp offer)	~80%
Rate card (“al-Ustath Mr”)	40%	Electricity, water, fines, traffic fines	60%
Rate card (“al-Ustath Mr”)	50%	Customs, zakat, tax	50%
Operator “Golden Mind” (DM)	50%	Fines (35,000 SAR settled for 17,500 SAR)	50%

The direct messages provided further evidence of the advertised funding source and payment methods. One operator described a trust-ramp starting near a 20% rate and stated that the business is money laundering and “all the loaders work this way”. Another required payment upfront in USDT-TRC20 and described the funding as “cards with balance that need draining”, a carding term, and subsequently described their own role as that of a broker. Settlement “proof” images showed large bills (in the hundreds of thousands of riyals) partly paid on a government portal, consistent with the fines and utility surfaces identified in the banking data.

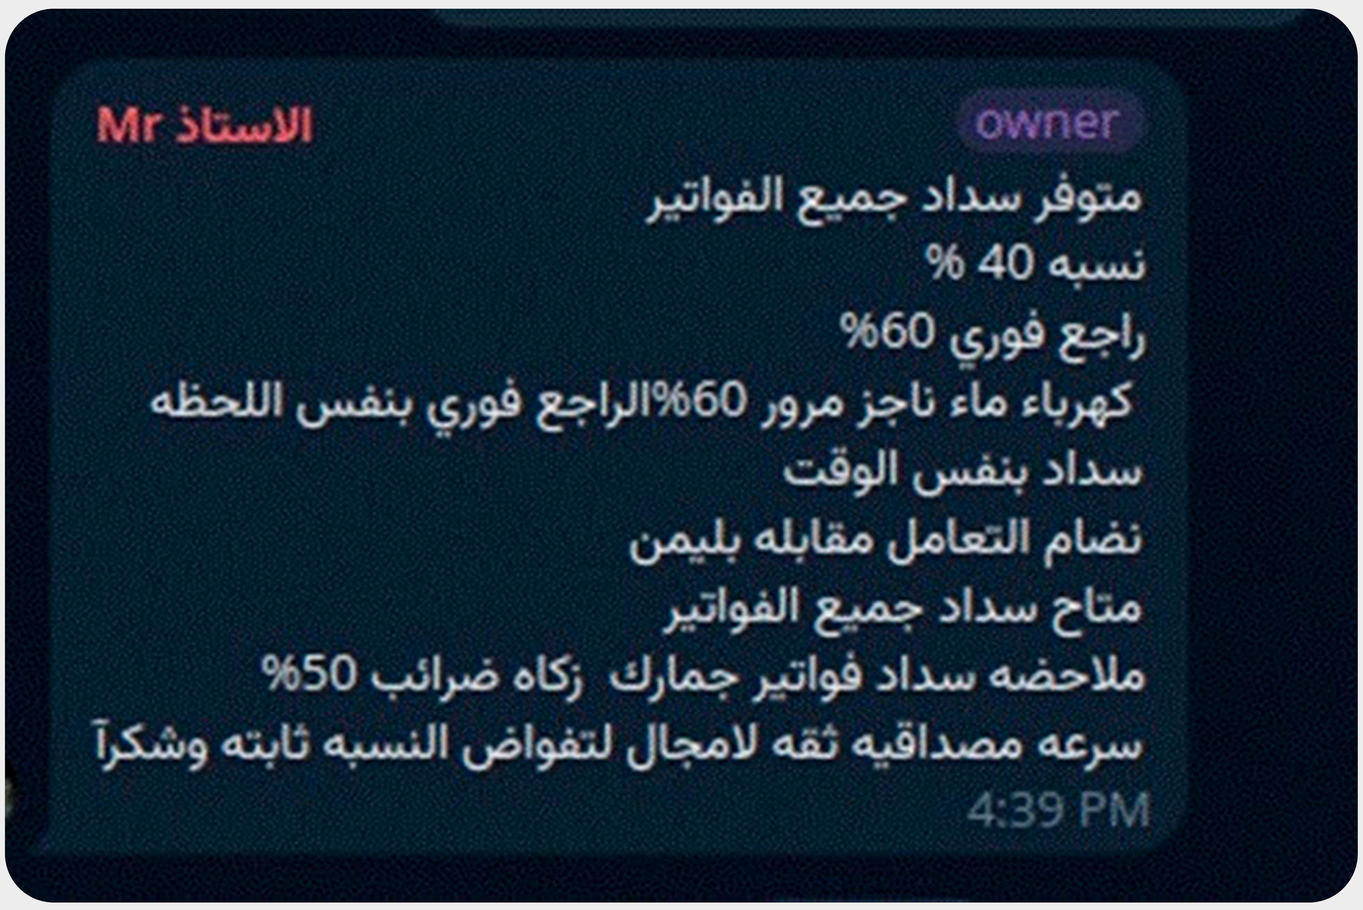


Figure 16. In-channel rate card from the operator "الاستاذ Mr".



Figure 17. DM thread with "Dark Baron" & "Golden Mind"

How the parties work together to cash out

Putting the roles and channels together, a single cash-out runs as follows:

- The operator advertises discounted fine or utility settlement to the public through the recruitment storefront.
- The customer supplies a biller reference or their national ID.
- The loader settles the bill on the government portal through the 3DS channel using Victim A's stolen card, with authentication completed through the compromised phone number or in-app approval following account takeover.
- The customer pays the operator the discounted amount in clean funds, and a broker or receiver may sit in the middle as escrow or as a mule.

Stage 7

Laundering and exit rails

The cash-out is itself the first laundering step: value leaves Victim A's account as an authenticated payment to a government entity, which can obscure the direct link to the fraud. The customer then pays the operator the discounted amount, by bank transfer, STC Pay, or USDT-TRC20, according to the payment methods advertised by the operators. Where TRC20 deposit addresses are available, on-chain analysis may help trace subsequent transfers and identify links to OTC services, other settlements or related operators, providing potential downstream investigative leads.

Why the scheme works

The scheme is profitable because it is relatively inexpensive to operate, can complicate tracing, and settles into a payment rail that is trusted by design. The operator's main upfront costs include a phishing kit and advertising spend. The advertised discounts may make the offer attractive to customers facing substantial government charges, particularly when the payment is presented as a legitimate settlement service.

Most importantly, the scheme is engineered to satisfy authentication controls using compromised credentials and access rather than simply bypassing them. Identity verification is satisfied through the compromised phone number or banking-app access, and GPS spoofing may help conceal the operator's location from banking-app controls. Viewed in isolation, these signals may not reveal the preceding compromise. Suspicion is displaced onto the two parties: the visible beneficiary is the customer, whose real bill was paid, and the loser is Victim A, whose account paid it, while the operator sits outside the government payment path.

Detection in practice

A single authenticated 3DS payment, seen in isolation, may appear legitimate, making the preceding account compromise harder to identify through transaction-level monitoring alone. Group-IB Fraud Protection's cross-channel visibility is what exposes it. The platform ingests session telemetry across three channels for the same banking identity: the web channel, the mobile-app channel, and the 3DS channel. This means an authenticated 3DS payment can be evaluated against the device, location and session history seen on the mobile and web channels moments earlier.

When a 3DS authorisation is satisfied, the platform checks the other channels for the same identity and associated device or session activity. In the confirmed cases, related mobile-channel activity had been flagged for indicators such as GPS spoofing, a roaming or hijacked eSIM, a newly registered device, or activity from a blocklisted location. A cross-channel correlation rule links the authenticated payment back to the takeover session that enabled it, which turns an individually “clean” transaction into a high-confidence fraud signal. Location-integrity and roaming/ eSIM signals dominated the confirmed cases, and those are exactly the signals that let the cross-channel logic connect an otherwise valid 3DS payment to a compromised session.

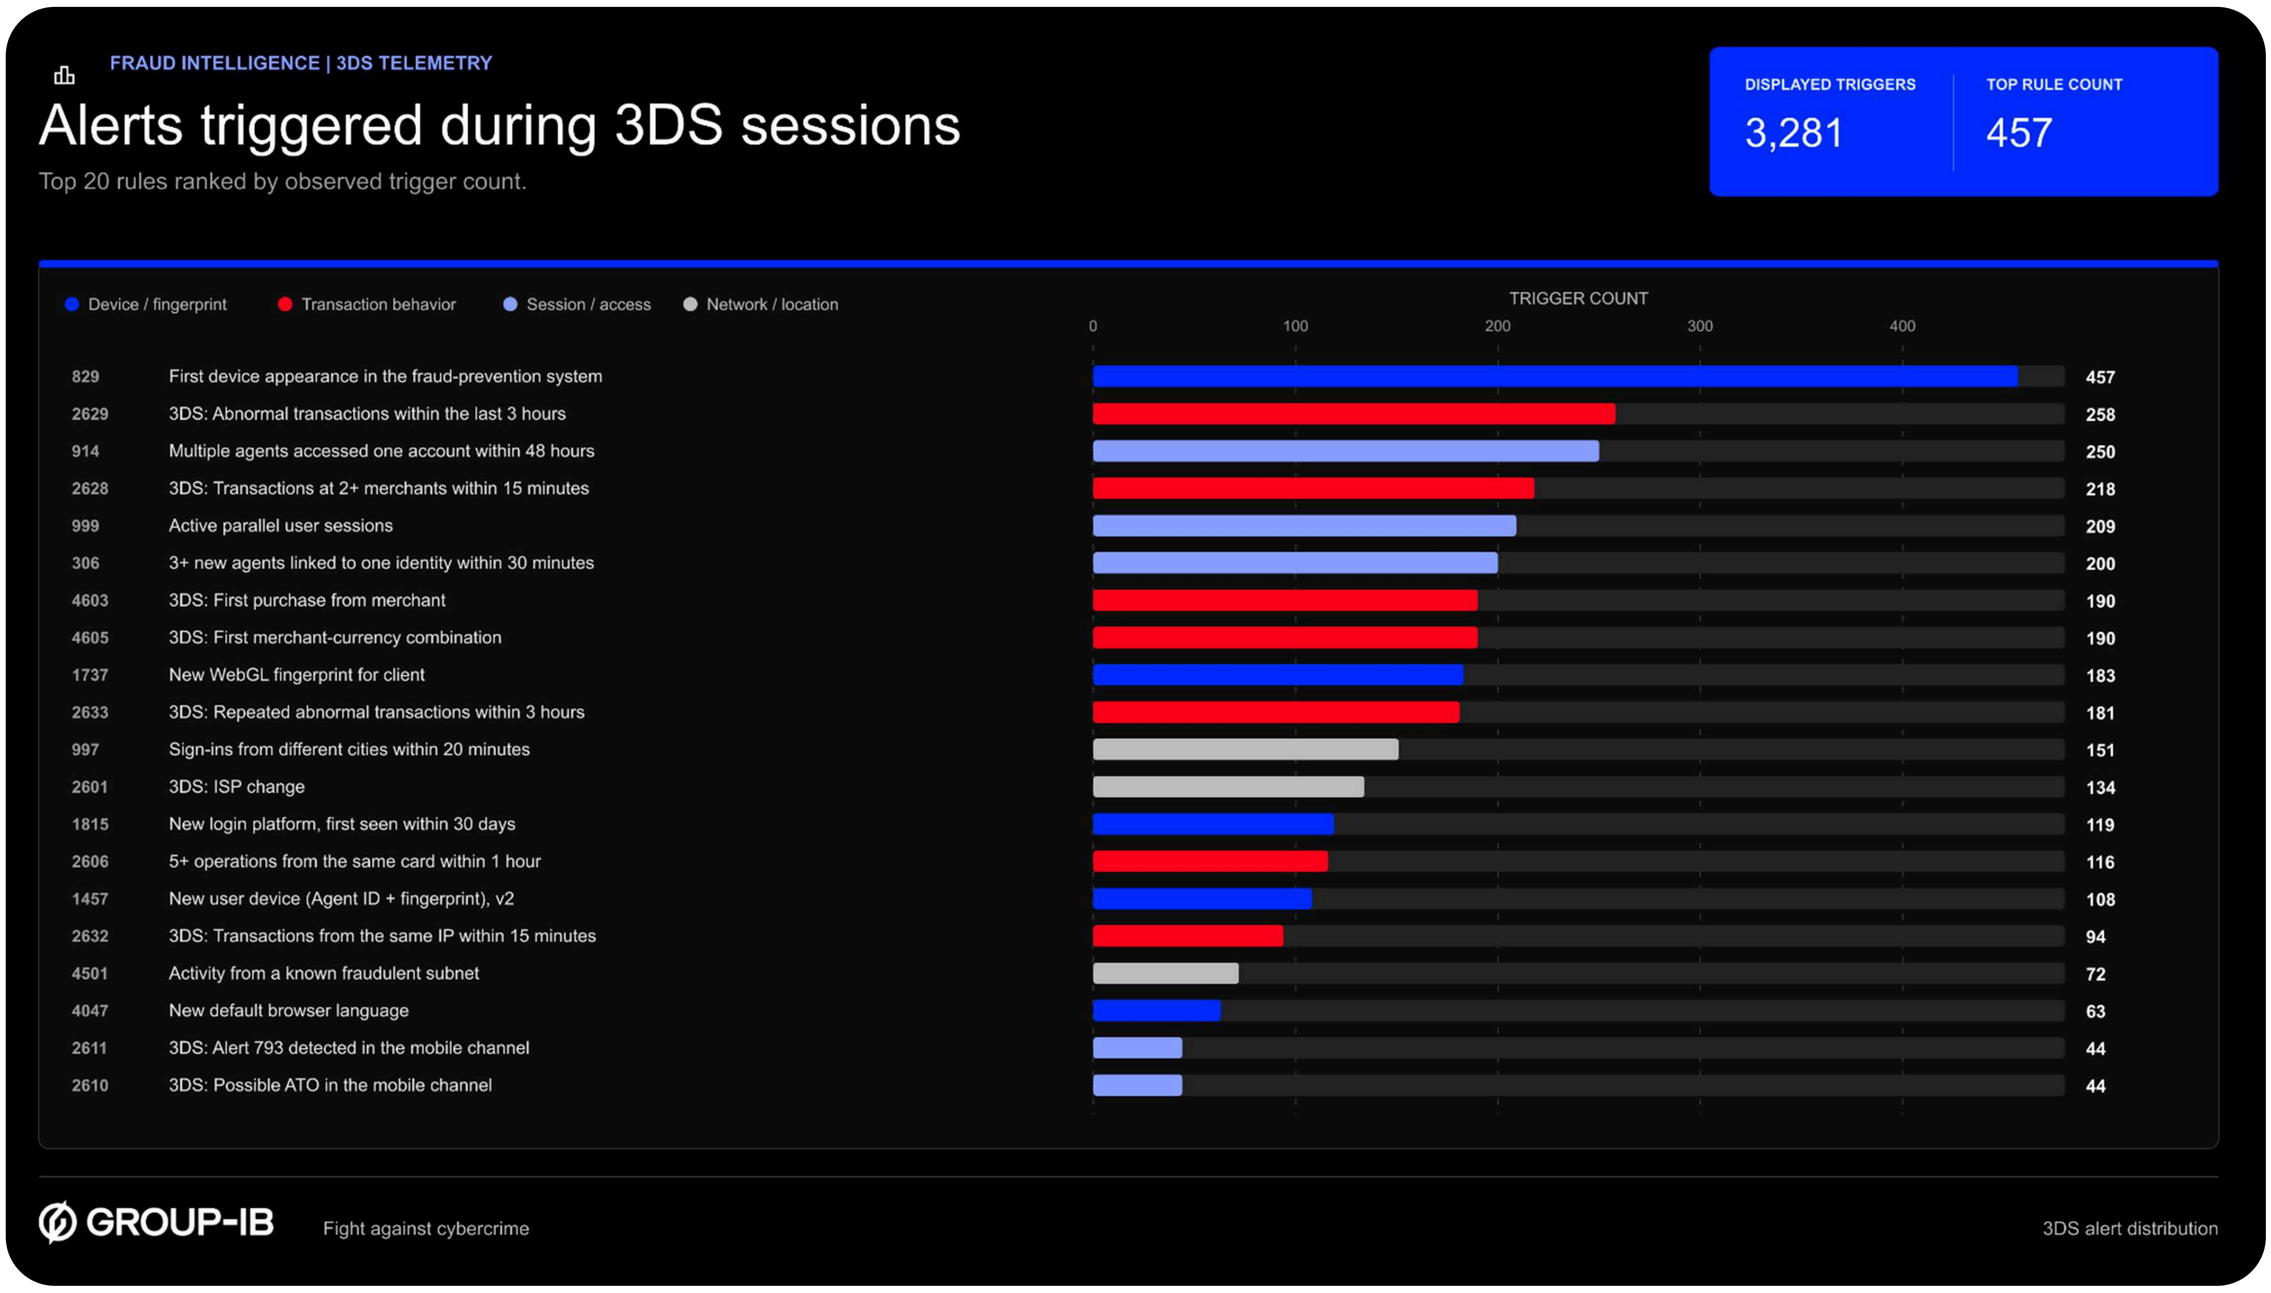


Figure 18. Detection signals observed in confirmed fraudulent 3DS sessions, grouped by theme, including cross-channel alerts indicating mobile-account compromise.

Impact

The confirmed impact, 80 compromised cards and USD 2.01 million across three institutions, represents the validated subset of the activity observed during the investigation. The wider reach of the scheme remains difficult to quantify because the model distributes harm across parties who each see only part of it.

Cardholders (Victim A)	Direct financial loss, account takeover, and the burden of discovering and disputing an authenticated transaction.
Paying customer	May be aware or unaware that the discounted settlement is funded by fraud. Customers who knowingly participate may face legal exposure, while those who are deceived may also suffer financial or legal consequences. The scheme can draw ordinary citizens into a criminal chain by presenting fraudulent payments as legitimate discounted services
Banks	Chargeback and liability exposure, and a detection problem that can be difficult to identify through transaction-level monitoring alone.
Government and portal operators	Potential erosion of trust in payment rails that depend on being seen as low-risk. The trusted rail is the very feature being abused.
Society	A scheme that turns two routine civic tasks, settling a fine and paying a utility bill, into a laundering channel, and can draw ordinary citizens into a criminal chain, whether knowingly or through deception.

07 Towards a predictive defence: Group-IB's cyber fraud fusion approach

The reason this scheme succeeds is structural: it is assembled from steps that each sit in a different defensive silo. Brand impersonation and malicious ads belong to digital risk teams, infrastructure and crypto flows to threat intelligence, the takeover and the 3DS payment to fraud teams, and the cross-institution money movement to no single institution at all. No single team necessarily has visibility across the entire chain, making coordinated detection and response essential. Cyber Fraud Fusion (CFF) addresses these gaps by coordinating five capabilities so that a detection in one domain can trigger action across the others. The scenario below illustrates how the fusion flow can be applied to this campaign and the three actor layers identified during the investigation.

Scenario: predicting and disrupting the bill-discount cash-out

Step 1

Detection (Digital Risk Protection)

DRP is the earliest detection point, and it targets the phishing layer. Monitoring paid search, domain registrations and hosting platforms surfaces one impersonation kit. Graph analysis then expands that single detection into the connected network: shared builder-account identifiers, shared nameservers and shared exfiltration accounts tie the 400-plus resources back to common operators. Automated, staged takedown can then be used to remove the fraudulent pages and intermediary booking redirectors, reducing the opportunity for further victim exposure.

Step 2

Intelligence enrichment (Threat Intelligence)

TI turns those detections into structured intelligence. It maps the exfiltration channels, the Jordan-registered infrastructure linked to the checker layer, and the Telegram cash-out ecosystem: the recruitment storefront, the processing channel, the shared @A_L_i_H_N contact, and the USDT-TRC20 addresses. Where transaction data is available, on-chain analysis can help trace relationships between wallets and settlements, while the combined infrastructure and behavioural intelligence can help profile actors and anticipate infrastructure rotation.

Step 3

Ecosystem alert and interception (Fraud Protection and CFIP)

Fraud Protection operationalises the intelligence into cross-channel rules: an authenticated 3DS payment to a government biller is re-scored against a simultaneous eSIM or roaming, GPS-spoofing and new-device signal on the same identity, which is exactly the correlation that exposed the confirmed cases. CFIP can then distribute anonymised, privacy-preserving signals across participating institutions, allowing a takeover pattern or other relevant indicator detected at one institution to inform risk decisions across the ecosystem.

Step 4

Attribution (Investigation)

Investigation traces customer payments and USDT-TRC20 flows through intermediary wallets towards cash-out points and, where legally accessible, seek KYC information associated with receiving accounts or exchanges. Combined with the Ramtha-linked device and infrastructure evidence, this can support an evidence package for law-enforcement referral and may help determine, using stronger identifiers, whether the three layers are operated by one group or several.

Outcome

Applied to this campaign, the fusion approach can create opportunities for coordinated intervention across multiple stages: phishing infrastructure can be identified and disrupted, cross-channel rules can increase the likelihood that an otherwise “valid 3DS” payment is detected, intelligence can be shared across participating institutions, and financial flows can be investigated further. A scheme designed to distribute activity across multiple domains can therefore become more visible and more effectively predicted, disrupted and attributed.

Why fusion matters

No single capability provides complete visibility into this scheme. DRP alone takes down sites the operator simply rotates, TI alone maps a network no one acts on, and a bank alone sees one authenticated payment and one isolated takeover. Together they form a defence that operates at the speed and scale of the attack: an architecture of defence that matches the architecture of deception.

08 Conclusion

This campaign is an illustration of a fraud that succeeds by exploiting the assumptions underlying security controls rather than necessarily breaking them. A geo-fenced phishing toolkit harvests a victim's identity, card data and phone number, the number is moved to an eSIM, the banking account is taken over, and a fully authenticated 3DS payment settles a paying customer's government bill. The cardholder is left facing the fraudulent charge, while the operator receives payment from a customer who may or may not be aware of the fraudulent source of the funds, turning a trusted government payment channel into part of the fraudsters' cash-out process.

The scheme is only consistently visible where the silos meet: in the correlation between an authenticated payment and the takeover that enabled it, and between the phishing front end, the underground marketplace and the crypto exit. That is precisely the seam a fused, cross-domain defence is built to watch. The presence of strong payment controls increases the complexity of the fraud chain, while a fused defence can turn that complexity against the operators by connecting signals that might otherwise appear legitimate or unrelated.

09 Recommendations

For affected institutions (banks and issuers)

- Treat account-recovery flows that rely on card data, PIN and SMS OTP as a potential point of compromise, and add device, behavioural and contextual risk checks to recovery and limit-change actions.
- Apply elevated scrutiny to high-value 3DS payments to government billers that follow a recent new-device login, an eSIM or roaming change, or a limit increase on the same identity.
- Where technically and legally feasible, coordinate with mobile operators to incorporate recent SIM or eSIM changes as risk signals during authentication.

For detection and monitoring teams

- Deploy cross-channel correlation across web, mobile and 3DS for a single identity, and re-score authenticated 3DS payments against concurrent takeover indicators rather than judging them in isolation.
- Prioritise location-integrity and roaming or eSIM signals (including impossible travel, GPS-spoofing indicators, new subnets), which were frequently observed across the confirmed cases
- Treat the government-biller destination as a potential cash-out indicator when it coincides with takeover signals, not as a low-risk one.

For government and portal operators

- Consider introducing additional risk signals for high-value or rapid repeat settlements, and support payer-side context so a stolen-card settlement can be distinguished more effectively.
- Provide banks and CERTs with a channel to flag portal transactions suspected of fraudulent settlement.

For customers and end users

- Access government and insurance services through official apps or bookmarks, not sponsored search results. A paid ad is not a trust signal.
- A steep discount on paying a government bill through an unknown third party may be a fraud or money-laundering lure. Using such services may expose customers to financial or legal consequences, particularly where they knowingly participate in suspicious transactions.

For the ecosystem

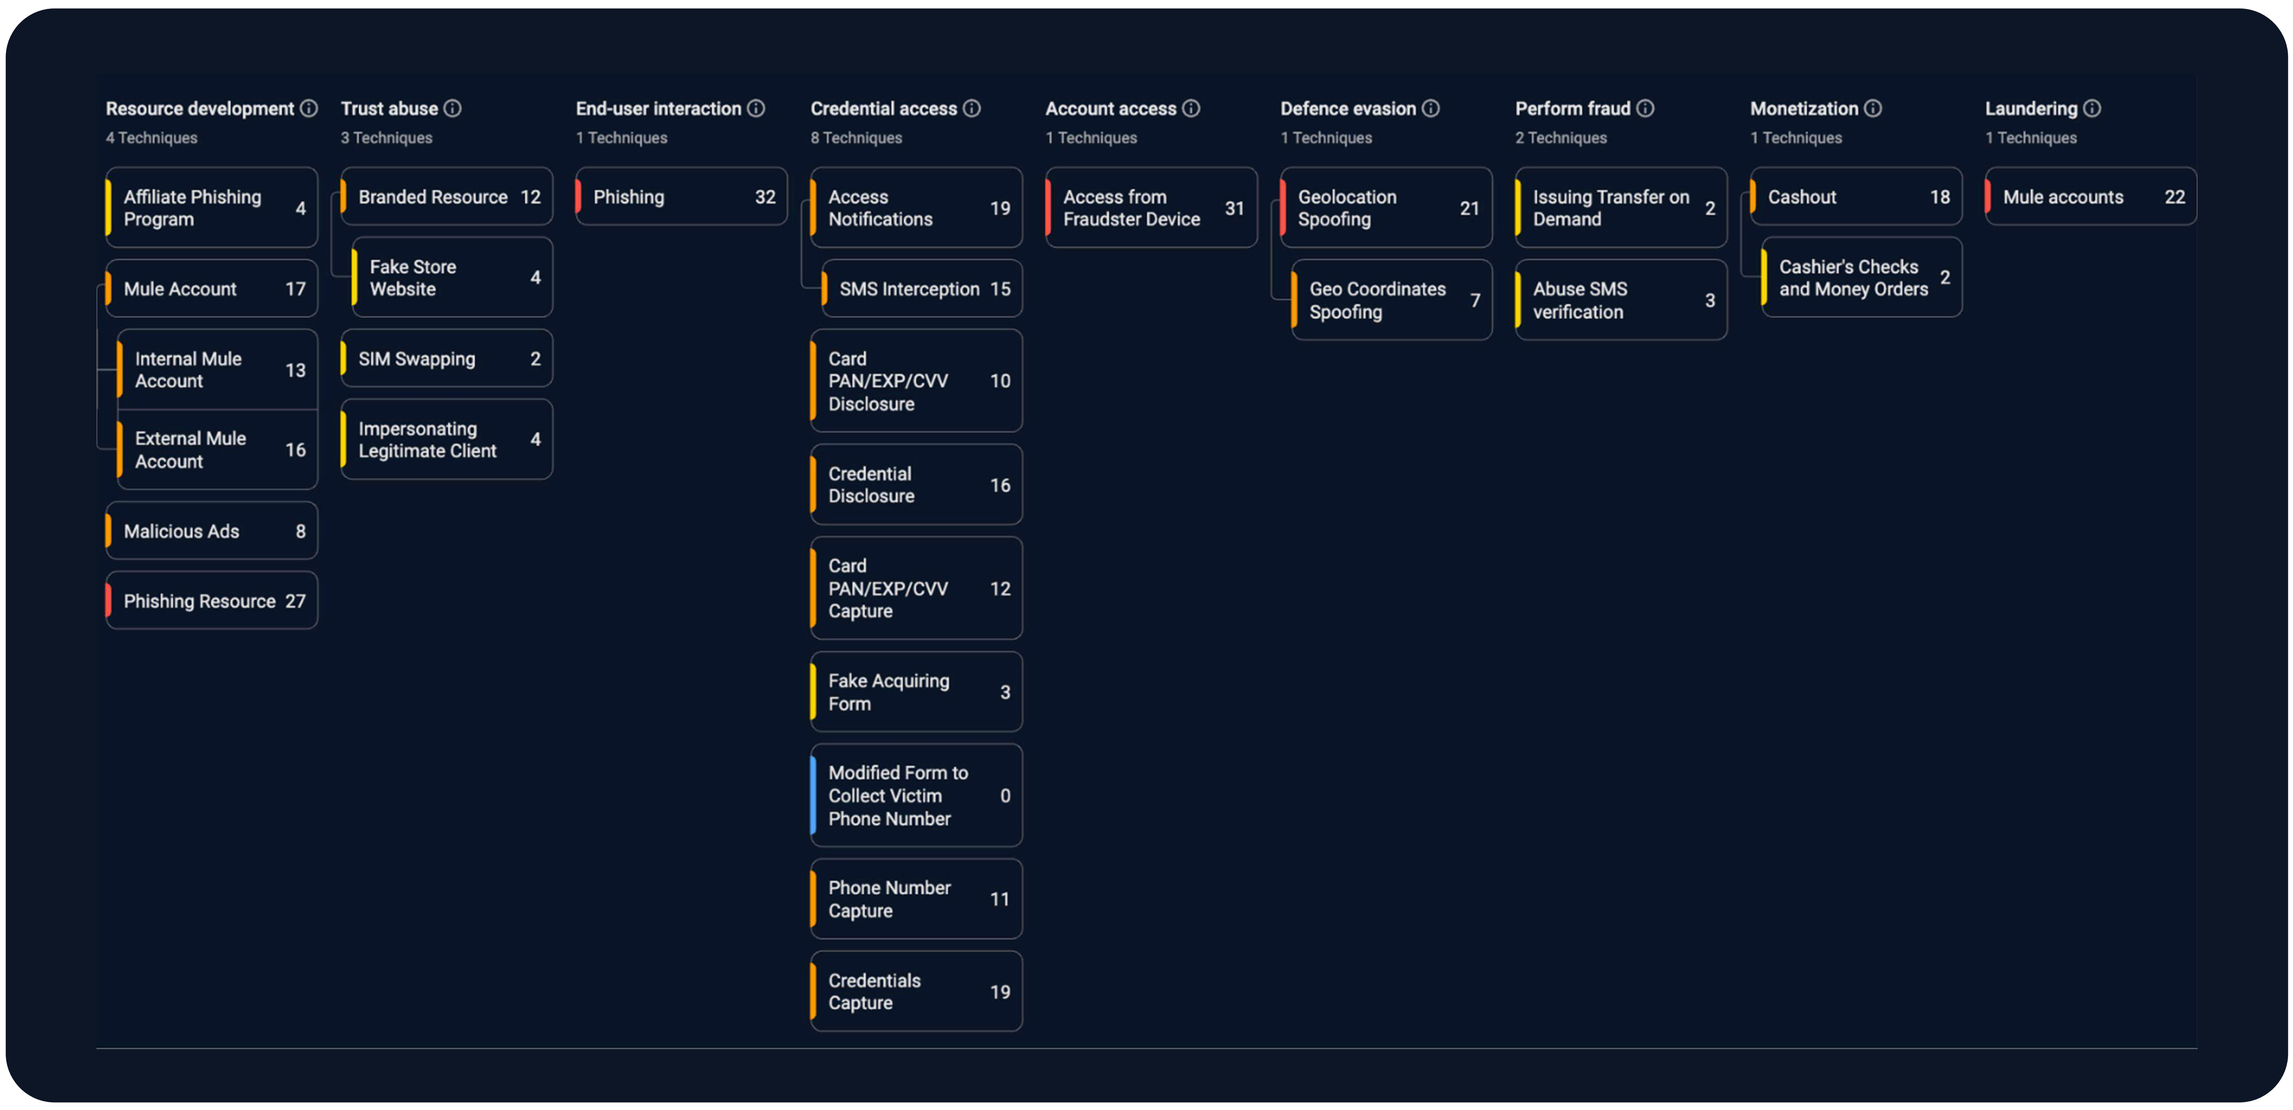
The scheme crosses domains, so the defence must too. Group-IB's Cyber Fraud Fusion combines Digital Risk Protection (impersonation detection and network takedown), Threat Intelligence (infrastructure and crypto-flow mapping), Fraud Protection and CFIP (cross-channel detection and real-time, privacy-preserving intelligence sharing across institutions), and Investigation (crypto tracing and attribution) into a coordinated approach designed to connect signals across the fraud chain.

10 MITRE ATT&CK

Indicative mapping of observed techniques, combining enterprise and mobile ATT&CK.

ID	Technique	How it appears here
T1583.008	Acquire Infrastructure: Malvertising	Paid search ads drive victims to phishing pages
T1583.001	Acquire Infrastructure: Domains	400+ rotating phishing domains across ~10 patterns
T1566	Phishing	Cloned government and insurance pages harvest data
T1598	Phishing for Information	PII, card, PIN and OTP collected via fake flows
T1451	SIM Card Swap (Mobile)	National identity approved eSIM issuance hijacks victim's number
T1078	Valid Accounts	Recovered banking credentials used as legitimate access
T1098	Account Manipulation	Password reset, limit increase, inter-account transfer
T1497	Virtualization/Sandbox Evasion	Geo-gating, CAPTCHA, bot-scoring, decoy shells
T1657	Financial Theft	3DS government-portal settlement and bill-discount cash-out

11 Fraud Matrix



12 Indicators Of Compromise (IOCs)

All indicators are defanged and provided for defensive use. . Several indicators below reference legitimate hosting, serverless and API providers whose services were abused by the operators without their knowledge or involvement; their inclusion is for defensive detection only and implies no wrongdoing by those providers. The following are partial IOCs. The full, current indicator set is maintained on the Group-IB Threat Intelligence portal and Fraud Intelligence portal.

01 Network indicators: phishing (examples by kit)

- Navigate TI or FI to see the detailed IOCs

02 Network indicators: exfiltration and control

- *[.]workers[.]dev: attacker-owned cloud-worker exfil endpoints (shared account)
- *[.]onrender[.]com: Realtime Socket.IO C2 collecting from multiple kits
- Supabase Realtime project backends (one per deployment)
- Commercial real-time messaging service used to operate the live operator panel (application key withheld pending confirmation that the infrastructure is inactive; available to customers via the TI portal)

- Telegram bot tokens embedded in kit source **(partially redacted in source data)**
- beacon.php: server-side PHP data and bot-gating endpoint (static kit)
- ipapi[.]co: third-party geo-check service abused for client-side gating

03 Underground and cash-out indicators

- Telegram recruitment storefront: customer recruitment
- Telegram 2D/3D processing channel: processing layer
- Visa-card carding supergroup (~1,400 members): gateway procurement
- Operator contacts: @A_L_i_H_N, @Mahmood100002 (as observed)
- USDT-TRC20 operator deposit address(es): captured, available via TI portal (not yet clustered on-chain)

04 Geohashes of real fraudsters geolocation

- svc95
- svc97

05 ISPs used by fraudsters

- Umniah
- Jordan Telecommunications PSC
- Jordanian mobile phone services
- Orange Jordan

Legal notice

Third-party trademarks. All product names, logos and brands referenced in this publication are the property of their respective owners. They are used solely for identification and descriptive purposes in the context of security research, and their use does not imply any affiliation with, or endorsement by, their owners. Screenshots of fraudulent resources are reproduced to document the impersonation of legitimate brands and services; the brands depicted are victims of that impersonation.

Third-party service providers. Certain indicators reference lawful hosting, serverless and API providers whose services were abused by the operators without their knowledge or involvement. No wrongdoing on their part is alleged or implied.

No legal advice. This publication is provided for information and defensive purposes only and does not constitute legal advice. Statements concerning the possible legal position of any party are general in nature; readers should obtain advice from a qualified adviser in the relevant jurisdiction.

Assessments and attribution. All assessments, attributions and confidence levels stated in this publication reflect the information available to Group-IB as at the date of publication and may change as further information becomes available. Where conduct is attributed to accounts, channels or groups, Group-IB makes no assertion as to the identity of any individual behind them beyond what is expressly stated.

GROUP-IB

Group-IB is a creator of cybersecurity technologies to investigate, prevent and fight digital crime.

1,600+
Successful investigations of high-tech crime cases

600+
Employees

60
Countries

$1 bln+
Saved by our client companies through our technologies

#1*
Incident Response Retainer vendor

*According to Cybersecurity Excellence Awards

11
Unique Digital Crime Resistance Centers

Global partnerships

- INTERPOL
- EUROPOL
- AFRIPOL

Recognized by top industry experts

- Forrester®
- datos INSIGHTS
- kuppingercole ANALYSTS
- Gartner®
- IDC
- FROST & SULLIVAN