Abnormal / Intelligence

THREAT INTELLIGENCE REPORT

Uncovering ZeroTokens:

Phishing Platform Uses Live Session Control to Harvest Financial Credentials and Identity Data for 50+ Financial Brands

Executive Summary

ZeroTokens is a human-operated phishing platform designed to collect financial credentials, identity information, and authentication factors through live, institution-specific workflows. Rather than sending every target through a fixed sequence, the platform gives operators real-time visibility into each active session and allows them to decide what appears next—from login prompts to one-time codes, identity checks, payment information, and transaction passwords. That human control is the platform's defining feature, turning a conventional phishing flow into an interactive collection process.

Abnormal AI first identified ZeroTokens through an Australian CommSec wave, but broader analysis revealed a much larger operation. From April 21 through September 25, 2026, ZeroTokens-linked activity delivered more than 1.2 million messages to 2,814 Abnormal customers and impersonated 25 banking and brokerage brands. Following a September 18 migration, the rebuilt console contained 58 brand-specific projects spanning multiple markets, up from 53 in the August 17 capture, showing that the CommSec wave was one deployment within a broader multi-brand system.

ZeroTokens is not a traditional automated adversary-in-the-middle platform. Its captured code does not connect directly to genuine financial institutions or steal authenticated session tokens. Instead, it is built to keep targets engaged while operators solicit app-push approvals and collect credentials, OTPs, email codes, identity data, and institution-specific authorization secrets in real time.

The operation has also shown significant resilience: when the original backend was placed under registry hold, the console was rebuilt on new infrastructure within roughly 20 hours. These findings show a scalable phishing operation whose strength comes less from a novel exploit than from coordinated human control, reusable infrastructure, and brand-specific workflows.

Attack Progression

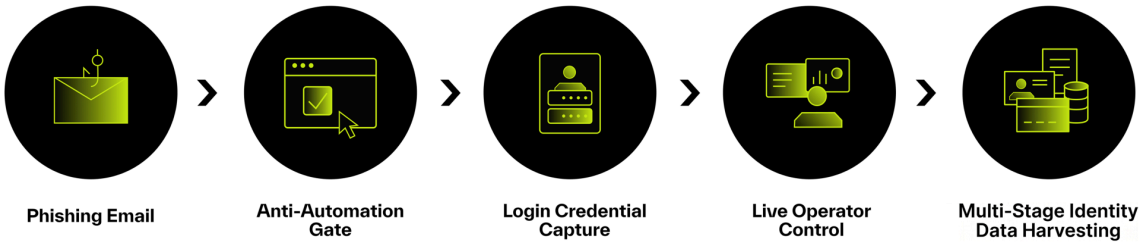

Phishing Email: Targets receive messages impersonating financial institutions and using account-related pretexts to create a plausible reason for action. Observed themes include W-8BEN recertification, account security, device enrollment, two-factor authentication activation, margin calls, and pre-IPO opportunities.

Anti-Automation Gate: Clicking the embedded link directs the target through a Cloudflare Turnstile gate designed to filter automated traffic before exposing the phishing flow. This helps reduce visibility to scanners and other automated analysis tools.

© 2026 Abnormal AI Inc. All rights reserved.

Login Credential Capture: Targets who pass the gate reach a cloned sign-in page tailored to the impersonated institution. The page captures account credentials while reproducing the branding and login options the target would expect from the legitimate service.

Live Operator Control: Once the target enters the phishing flow, their session appears in ZeroTokens' live session board. A human operator can monitor submitted information and decide which preconfigured screen or response the target sees next, including retry prompts when a verification attempt fails.

Multi-Stage Identity Data Harvesting: The operator can guide the target through additional prompts for one-time codes, app approvals, identity information, and institution-specific authorization credentials. The exact sequence can vary by project and by how the target responds during the session.

Impact Analysis

A single phishing session can expose far more than login credentials. ZeroTokens can move targets through institution-specific prompts for authentication responses, identity information, and authorization credentials. A successful interaction can therefore create multiple forms of financial and identity exposure that a password reset alone may not address.

Not all MFA provides the same protection against live phishing. ZeroTokens is designed to collect one-time passcodes and email codes and solicit app-push approvals while the target remains engaged. These user-transferable factors can be relayed during an active session, while phishing-resistant methods such as FIDO2/WebAuthn security keys and passkeys are not susceptible to the same mechanism.

Brand-specific workflows make the platform reusable across markets. The project library expanded from 53 to 58 financial brands, while observed phishing activity impersonated 25 brands. That combination allows the same operator-controlled model to scale across institutions without relying on a single generic phishing flow.

Individual infrastructure takedowns have limited impact on the broader operation. ZeroTokens rotates disposable sender and lure infrastructure and rebuilt its control plane on new infrastructure within roughly 20 hours after the original backend was placed under registry hold. Defenders should therefore prioritize recurring platform behaviors and operator-controlled infrastructure rather than isolated domains alone.

Strategic Actions for CISOs

Prioritize Phishing-Resistant Authentication: Move high-risk users and financial workflows away from authentication methods that depend on transferable codes or user-approved prompts. FIDO2/WebAuthn security keys and passkeys prevent the type of real-time relay ZeroTokens is designed to exploit because the authentication challenge is bound to the legitimate origin.

© 2026 Abnormal AI Inc. All rights reserved.

Expand Incident Response Beyond Password Resets: Reconstruct the full phishing interaction, including which screens the target saw and what information or authentication responses they provided. Treat any OTP or email code entered during the live interaction as potentially used in real time, and investigate any app-push approval the target was convinced to authorize. Rotate exposed institution-specific authorization credentials and coordinate with the affected financial institution when appropriate.

Account for Financial and Identity Data Exposure: Include fraud and identity teams when responding to suspected ZeroTokens compromise. The platform can collect identity information, card data, transaction passwords, and other institution-specific details alongside login credentials. Warn targets who reached the pre-IPO stage about the possibility of related follow-on outreach.

Strengthen Detection Beyond Sender Authentication: Do not treat SPF, DKIM, or DMARC alignment as evidence that a message is benign. Hunt for recurring sender, email service provider, lure fleet, and multi-stage collection patterns that persist even as individual domains change.

Prioritize Durable Platform Infrastructure: Focus blocking, hunting, and provider abuse reporting on confirmed ZeroTokens control plane infrastructure, operator-controlled domains, account identifiers, and recurring platform signatures. Avoid broad blocking of shared hosting, CDN, or email provider infrastructure that could affect legitimate services.

© 2026 Abnormal AI Inc. All rights reserved.

Attack Overview

ZeroTokens is built around a common attack model in which targets are directed to institution-specific phishing sites and guided through a series of preconfigured collection and verification stages under live operator control. The exact brand, lure, authentication method, and information requested can vary by deployment, allowing the flow to be tailored to the financial institution being impersonated.

The example below shows how that model operates in the CommSec deployment that first surfaced the ZeroTokens platform. The attack impersonated CommSec—an Australian brokerage service within the Commonwealth Bank group—using W-8BEN recertification emails and a Cloudflare Turnstile gate. Targets were led into an operator-controlled sequence designed to collect login credentials, authentication responses, identity information, and authorization credentials. Later campaign-wide analysis showed this deployment was one implementation of a much broader operation spanning multiple financial brands and markets.

Email Delivery Infrastructure

The Australian CommSec wave analyzed here used ten sender domains—four `.com` and six `.asia`—all registered through one registrar, on one nameserver set, and aged five to eleven months before use. None published a mail exchange (MX) record, meaning the domains could send email but were not configured to receive it, and none had been captured by the Wayback Machine. Researchers also observed twelve submission HELO strings, all of which were forged. Three did not resolve, while the other nine belonged to unrelated third parties.

Messages in the Australian CommSec wave were sent through nine abused SendGrid accounts, and the sender domains had no prior sending history. The messages carried SendGrid's DKIM signature and passed SPF, DKIM, and DMARC checks, helping the activity avoid obvious email authentication-related red flags. Each of the ten sender domains was tied to a single subject template, with no rotation between domains.

These domains and SendGrid accounts represent the delivery infrastructure observed in the CommSec wave, not the full ZeroTokens-linked sender population. Campaign-wide analysis identified additional delivery patterns across the five-month operation, including earlier infrastructure generations that predated the August CommSec activity. Those broader patterns are detailed later in **Campaign Scope and Evolution**.

Phishing Email

In the CommSec example, the phishing email used W-8BEN tax documentation recertification as its pretext. Non-US individuals use Form W-8BEN to certify their foreign status for US tax withholding and reporting, making it relevant to customers who hold US investments through a brokerage. That gives the lure a plausible reason to ask the target to review account information and helps make the identity and authorization requests that follow appear consistent with the stated purpose. W-8BEN recertification was also the most common observed pretext across the broader ZeroTokens-linked campaign, although the operation used several other financial and account-security themes.

© 2026 Abnormal AI Inc. All rights reserved.

The first W-8BEN email template (Fig. 1) tells the recipient that their W-8BEN tax documentation requires review to keep account records current and directs them to an “official account service portal” to complete the update.

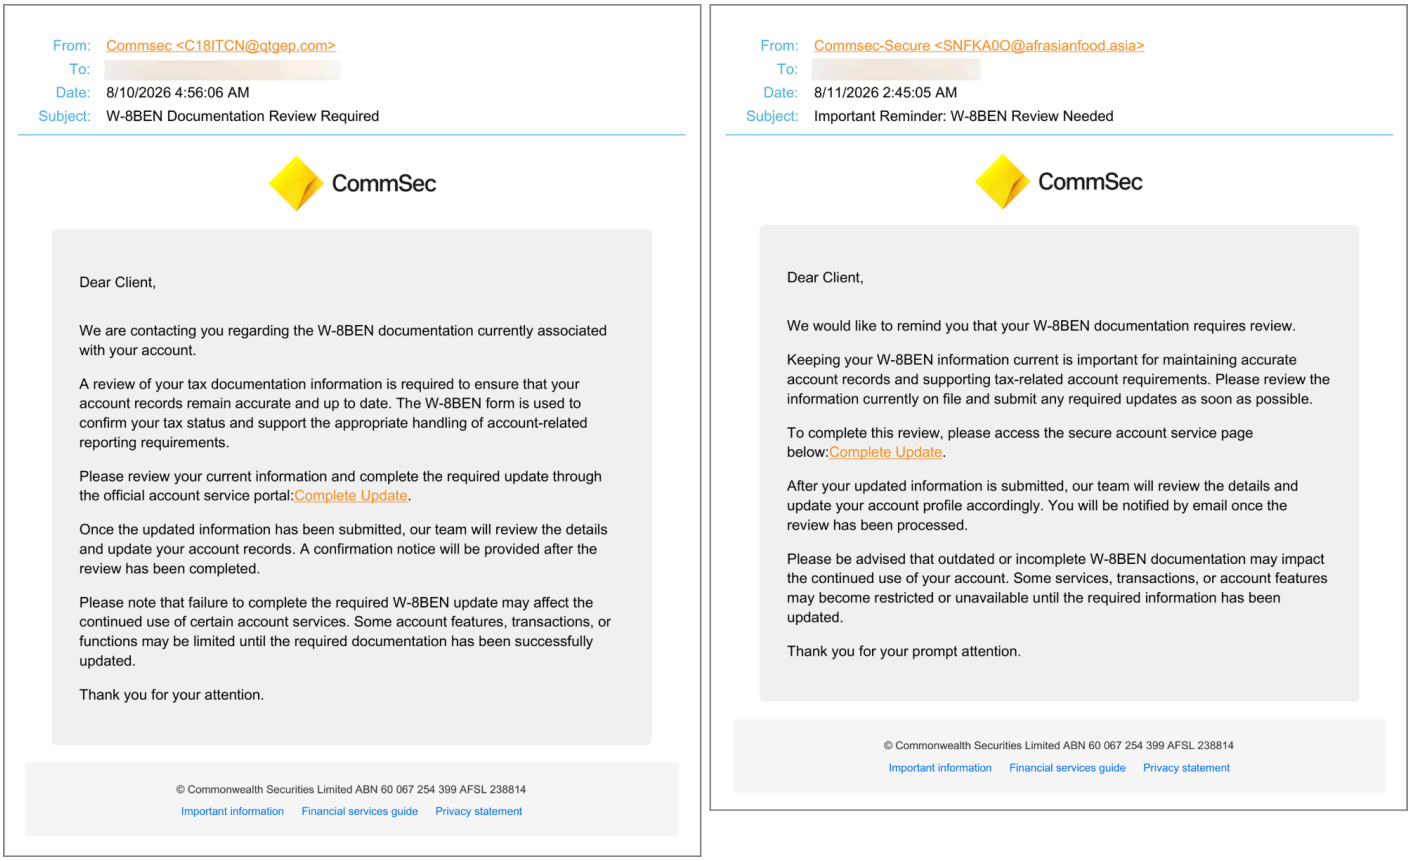

Fig. 1: The display name impersonates the broker, while the underlying sender address uses a random local part on a throwaway domain.

The message avoids threatening immediate suspension, warning instead that some services, transactions, or account functions may be limited if the review is not completed, giving the lure a restrained, administrative tone.

The second template (Fig. 1) uses the same underlying pretext but presents it as a reminder, asking the recipient to review and update their W-8BEN information. It similarly avoids overt urgency, warning only that outdated or incomplete documentation may eventually restrict certain account features or transactions—language that helps the message resemble a routine compliance notice instead of an obvious phishing lure.

Anti-Automation Gate

Before reaching the phishing flow, the target encounters a Cloudflare Turnstile gate designed to filter automated traffic (Fig. 2). The gate functions as an anti-automation measure rather than a security boundary; the phishing content remains part of the same deployment even though automated tools may have difficulty reaching it.

© 2026 Abnormal AI Inc. All rights reserved.

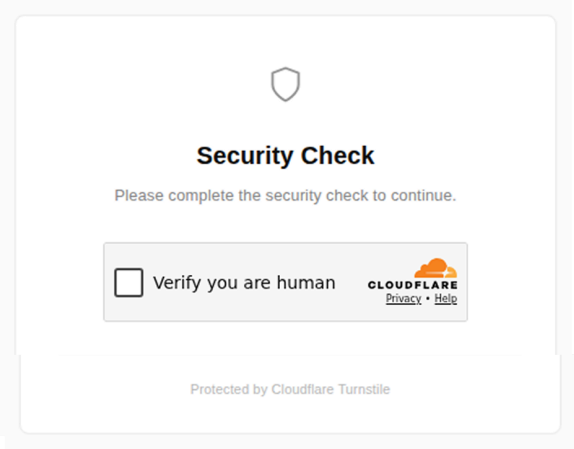

Fig. 2: The Turnstile gate retrieves its sitekey at runtime rather than embedding it in the page.

The Turnstile sitekey is retrieved from `api.0tokens[.]xyz/turnstile/site-key/` at request time, allowing the gate configuration to be changed centrally across deployments.

Multi-Stage Collection Flow

Note: The numbered steps below describe a representative CommSec attack flow and do not correspond to the step numbers embedded in the phishing pages themselves. Those breadcrumbs are part of the target-facing templates; because ZeroTokens lets the operator move between available screens dynamically, they do not necessarily reflect the order in which the interaction unfolds.

Step 1: Credential Capture and Live Session Control

Behind the gate, the target lands on a high-fidelity imitation of the selected financial institution. In the CommSec deployment analyzed here, the site reproduces the broker's branding along with legitimate Australian Business Number (ABN) and Australian Financial Services Licence (AFSL) information, genuine privacy language, and functional help links. Within that imitation, ZeroTokens can display up to eight different target-facing screens modeled on stages of the institution's verification process.

The flow begins with a cloned sign-in page designed to capture the target's account credentials (Fig. 3). Because CommBank customers can access CommSec using their NetBank credentials, the phishing site mirrors both login methods. Targets first see the standard CommSec credential form on the left. They can either submit their CommSec credentials and continue through the attack flow or click the "Login with NetBank" button, which loads a separate NetBank credential capture page before continuing through the same flow.

© 2026 Abnormal AI Inc. All rights reserved.

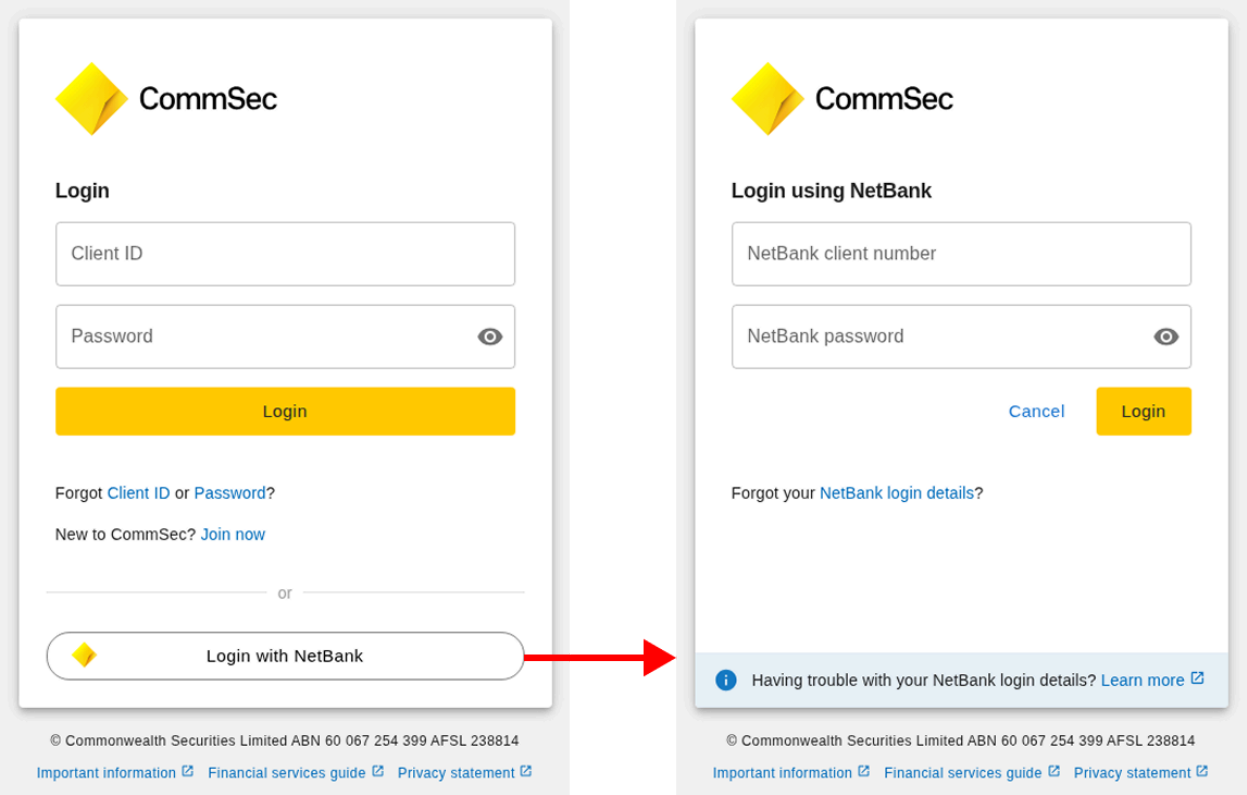

Fig. 3: The initial CommSec login page (left) offers a "Login with NetBank" option that loads a separate NetBank credential capture page (right).

These two CommSec login options illustrate how closely the clone tracks the genuine institution's sign-in experience. That same institution-specific design extends across a much broader project library spanning multiple brands and markets. The console build captured on August 17 contained 53 institution-specific projects across 14 locales, including configurations for US brands such as Fidelity, E*TRADE, Merrill Lynch, Public.com, TradeZero, Avenue Securities, PayPal, comdirect in Germany, Rakuten Securities in Japan, and Kiwoom Securities in South Korea.

A project's presence in the platform does not necessarily establish that it was used in an active or previously active campaign. It does, however, show that ZeroTokens was designed to support institution-specific phishing workflows across a broader set of financial brands and markets than those confirmed in observed activity.

Across those projects, the screens and verification steps may vary, but the underlying live-control model remains consistent. As the target enters information into either version of the CommSec login page, the field values are reported back to ZeroTokens. Those values appear in the live session board alongside the target's current status, giving the operator visibility into the interaction as it unfolds (Fig. 4).

© 2026 Abnormal AI Inc. All rights reserved.

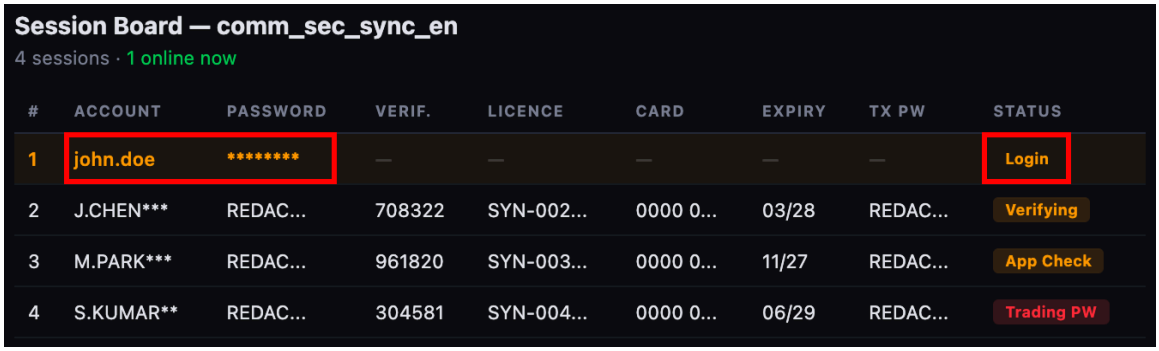

Fig. 4: Simplified reconstruction of the live session board, which consolidates active targets and collected information into a single operator view.[1]

Rather than advancing every target through the same fixed sequence, ZeroTokens lets the operator use that visibility to control the flow in real time. From the same interface, the operator can open a per-session command menu and select which screen to present next based on how the interaction is progressing (Fig. 5).

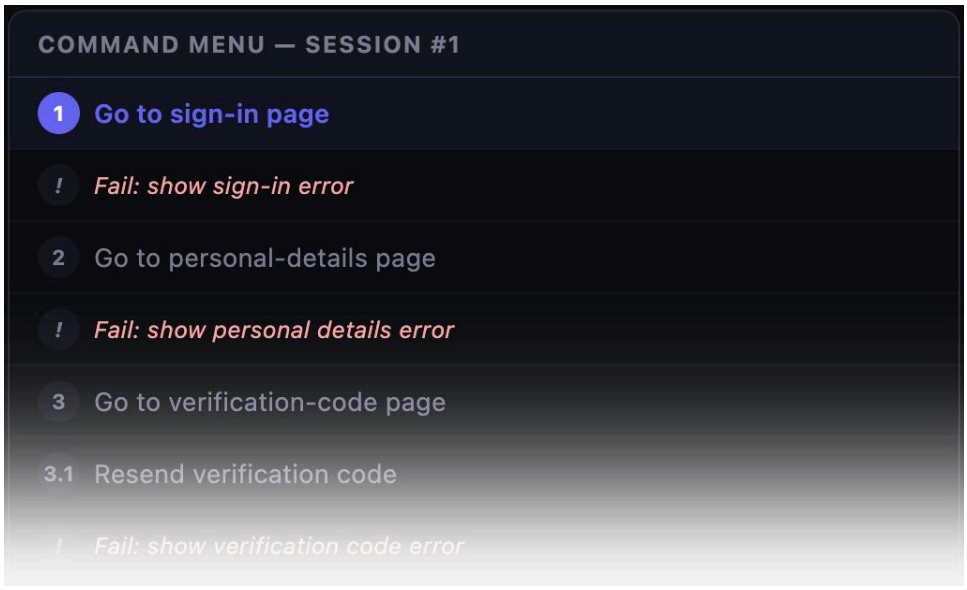

Fig. 5: A session-specific command menu exposes the available branches of the phishing workflow.

With the initial credentials collected, the interaction can move beyond account sign-in into additional verification and information-gathering stages.

Step 2: MFA Collection

ZeroTokens is designed to extend phishing beyond initial sign-in by reproducing the authentication and verification steps used by different financial institutions. Those steps vary by project, with captured configurations supporting one-time codes, app-based approvals, email verification, and institution-specific mechanisms such as mobileTAN and photoTAN. This lets operators adapt the flow to the target's authentication method, selecting the appropriate branch and presenting the corresponding screen.

[1] *Populated values are synthetic; subsequent session board figures use the same treatment.*

© 2026 Abnormal AI Inc. All rights reserved.

In the CommSec project, this capability takes the form of two MFA branches modeled on CommBank authentication: entering a six-digit SMS verification code or approving a login through the CommBank app.

For code-based verification, the operator can select “Go to verification-code page,” which serves a CommSec-branded screen asking the target to enter a six-digit SMS security code (Fig. 6). The field accepts up to six characters and uses `autocomplete=one-time-code`, allowing supported browsers and devices to recognize it as an OTP field and potentially offer the code for autofill. This enhances the appearance of legitimacy and reduces friction.

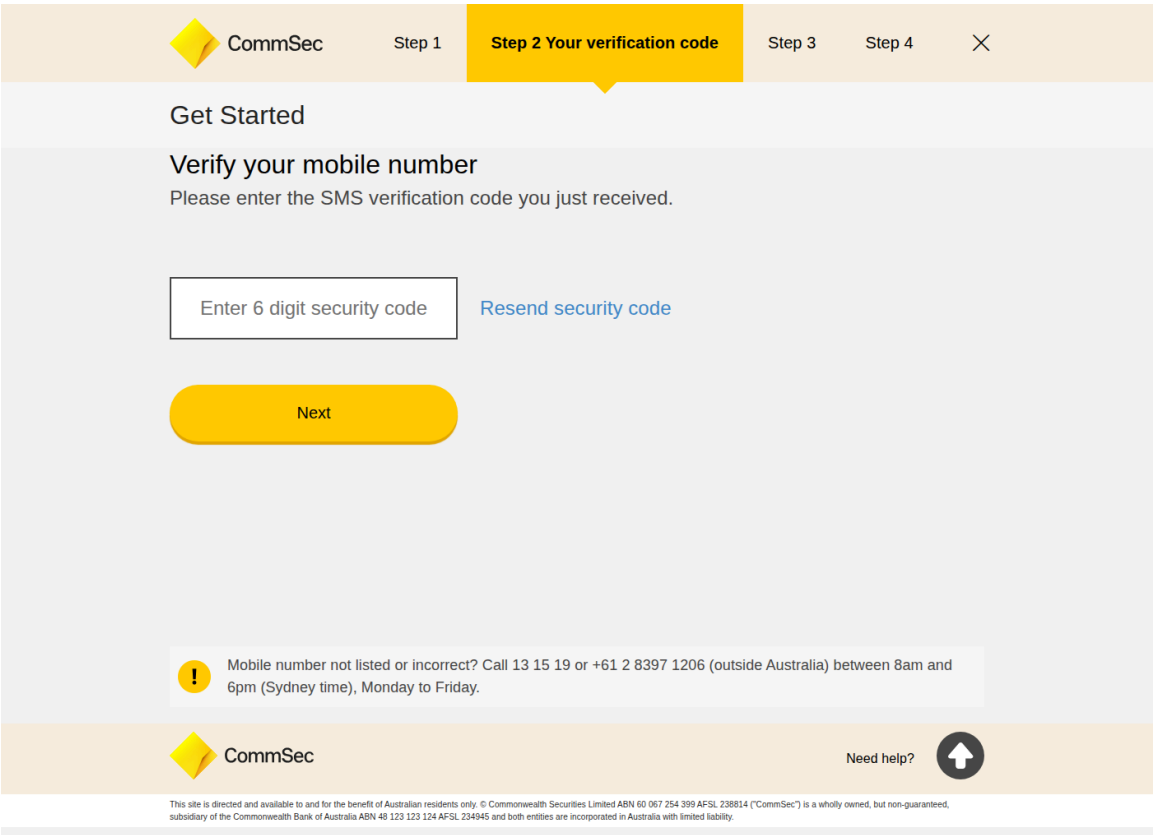

Fig. 6: The SMS branch embeds the six-digit verification-code prompt within the same staged CommSec interface used throughout the phishing flow.

Alternatively, the operator can select “Go to app auth page,” which presents a CommBank-branded screen instructing the target to approve a login through the CommBank app (Fig. 7). This path does not ask the target to relay a code; it attempts to induce an authentication action on the target’s own device.

© 2026 Abnormal AI Inc. All rights reserved.

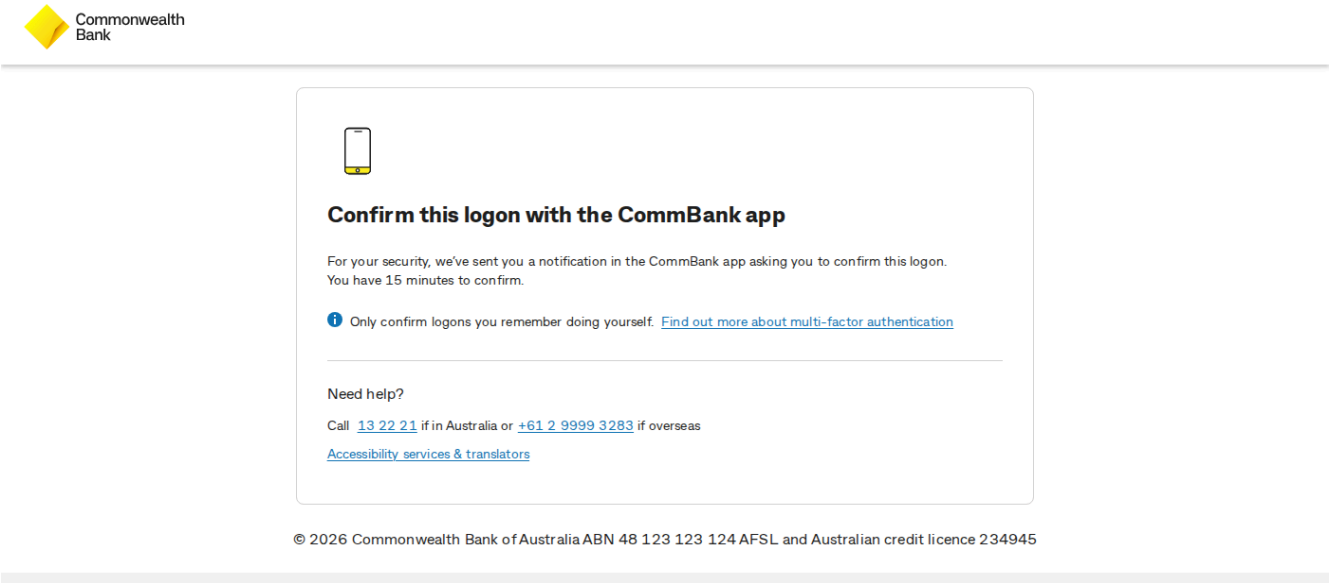

Fig. 7: The app-auth branch shifts the interaction from entering information into the phishing page to approving a login through the target's banking app.

These screens are designed to support an authentication attempt outside ZeroTokens. The platform does not connect to the genuine institution or generate an MFA challenge on its own. However, an operator could use credentials captured in the previous step to attempt a separate login at the real institution, then present the corresponding ZeroTokens screen if that login triggers an SMS code or app-approval request. Whether operators used the harvested credentials this way during the observed campaign was not directly confirmed.

The session board also reflects these authentication stages. A dedicated verification field appears alongside the captured account and password, while session states distinguish targets undergoing code verification from those at the app-approval stage (Fig. 8).

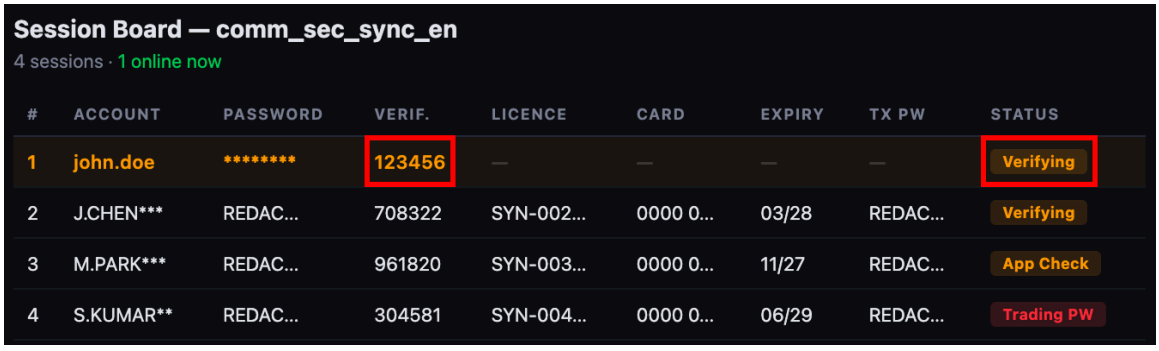

Fig. 8: The session board includes a dedicated verification code field and separate MFA-stage statuses.

From there, the operator can continue the same live interaction into screens that request additional identity and account information, extending the collection beyond credentials and authentication.

© 2026 Abnormal AI Inc. All rights reserved.

Step 3: Identity Data Collection

The next step in the attack flow is determined by the project configuration. Each financial institution project has a harvest schema that defines the details ZeroTokens is designed to collect. These vary with the institution being impersonated, ranging from identity and card information to specialized authentication and account data.

In the CommSec project, the operator can select “Go to personal-details page,” which presents a screen requesting the target’s driver’s license details (Fig. 9).

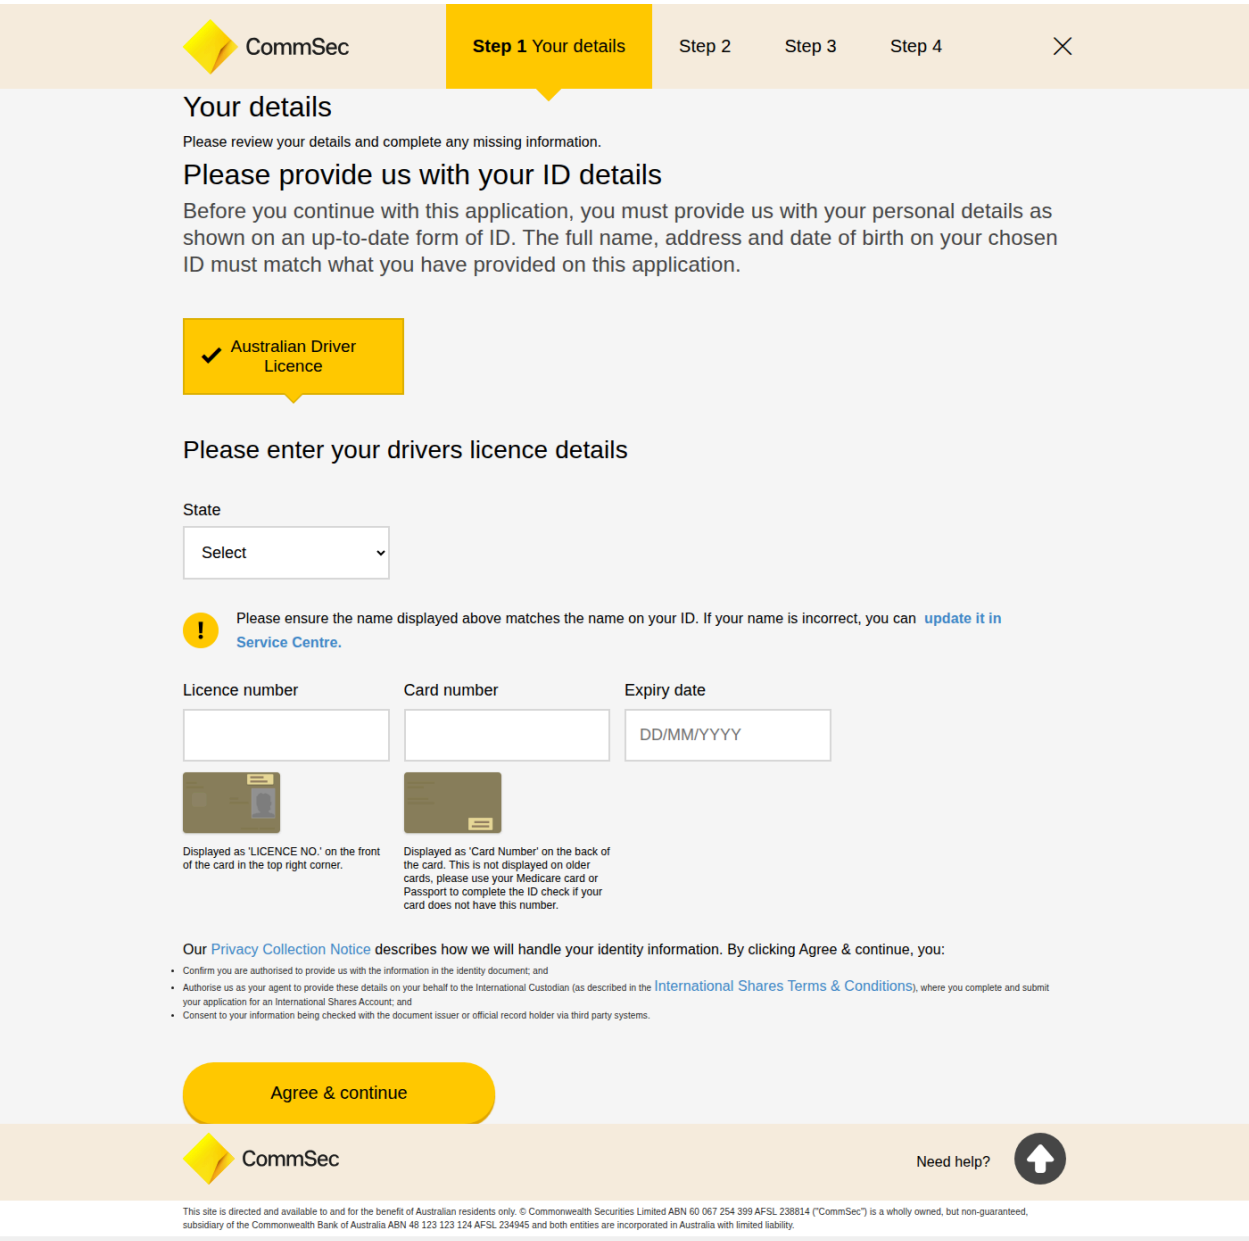

Fig. 9: The fraudulent CommSec verification page requests the target’s driver’s license details.

Once submitted, those values populate dedicated fields in the same session record that already contains the target’s account, password, and verification data (Fig. 10).

© 2026 Abnormal AI Inc. All rights reserved.

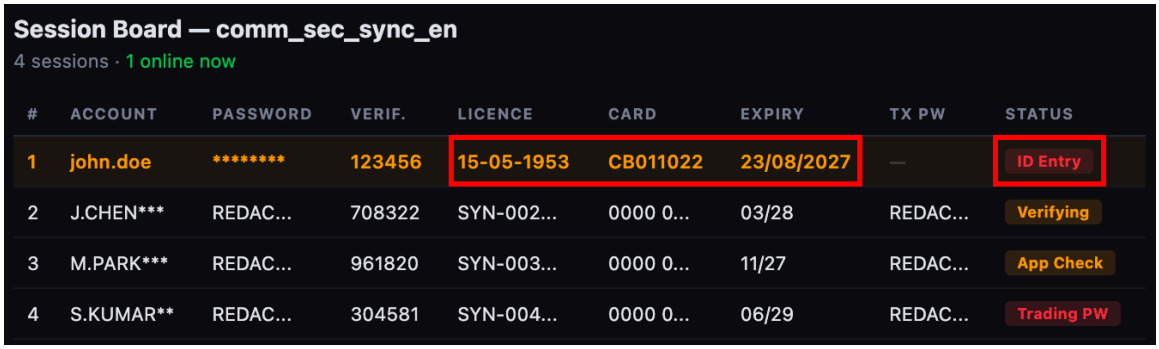

Fig. 10: The CommSec session schema places identity and card fields alongside information collected during earlier stages of the interaction.

By this stage, a single ZeroTokens session can contain account credentials, authentication data, identity information, and card details collected across successive interactions. That breadth of collection is a defining feature of the platform: rather than ending once access credentials are captured, the workflow can continue building a richer record of the target.

Step 4: Trading Password Collection

ZeroTokens' institution-specific schemas can also include specialized credentials and authorization data tied to how individual financial institutions operate. Captured projects include cloud-certificate fields and a 12-digit authorization code for Kiwoom Securities, mobileTAN and photoTAN fields for comdirect, and an emoji-grid selector for Rakuten Securities.

In the CommSec project, that specialized credential is a Trading Password, which is used to authorize orders rather than to sign in to the account. The operator can select "Go to payment-password page" and present the target with a CommSec-branded screen that requests the Trading Password (Fig. 11).

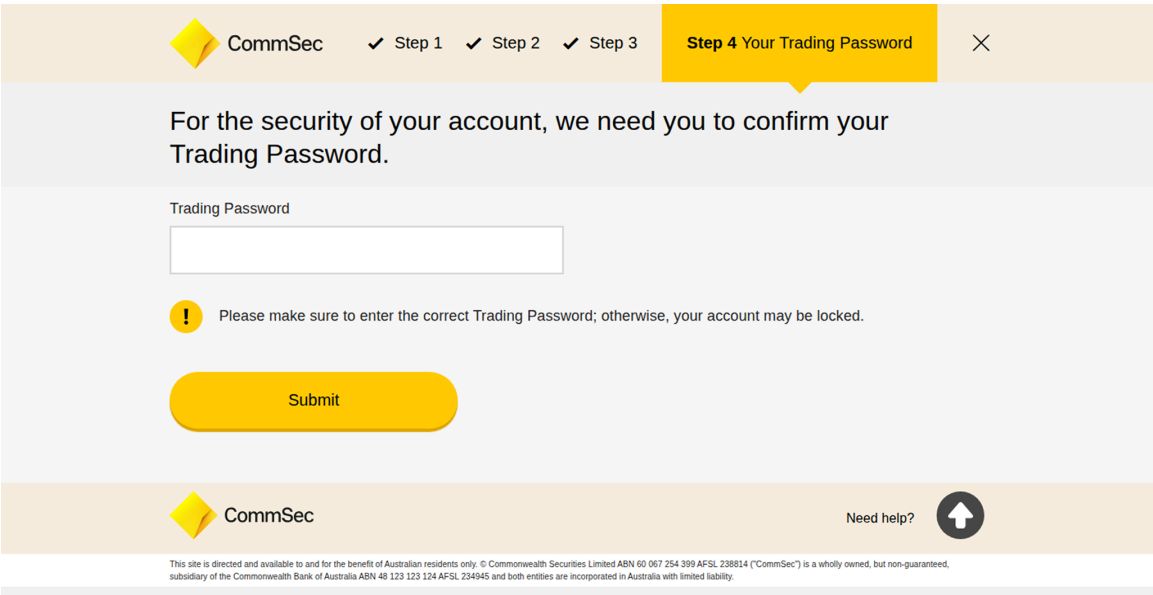

Fig. 11: The phishing page frames the Trading Password request as an account security confirmation.

© 2026 Abnormal AI Inc. All rights reserved.

The significance of this stage is not simply that ZeroTokens collects one more piece of data. It's that it expands the potential exposure from account authentication to a separate authorization credential that may be relevant to downstream financial activity.

Driver's license details are consistent with identity verification or account-recovery processes, while a trading password can be used to authorize certain account actions. ZeroTokens contains no functionality to perform those actions itself, but the field set could provide information relevant to higher-risk activity conducted outside the platform.

Fig. 12: The highlighted session shows account credentials, a verification code, license and card details, and a trading password consolidated into a single record.

In the CommSec flow, the trading password request marks the last collection stage before the interaction moves to its terminal screen.

Step 5: Terminal Screen and Potential Follow-On Fraud

The CommSec flow ends with an unexpected change in narrative. After beginning with a W-8BEN tax recertification lure and progressing through account verification screens, the final page tells the target that an application for a SpaceX pre-IPO share allocation has been submitted (Fig. 13). No pre-IPO theme appears in the campaign's 10 observed email templates, the preceding target-facing screens, or the information ZeroTokens is configured to collect.

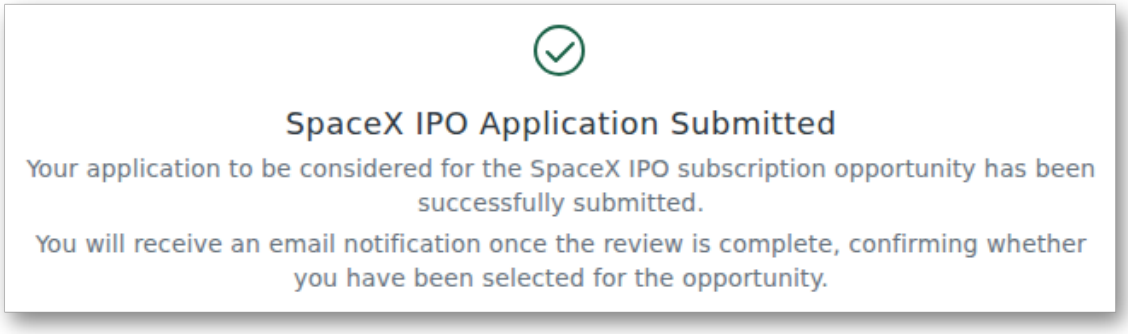

Fig. 13: The final screen introduces a SpaceX pre-IPO subscription narrative not seen elsewhere in the observed CommSec flow.

The disconnect suggests the page may be intended to bridge into a separate investment fraud interaction—for example, a later message claiming the target received an allocation and requesting payment. It could also indicate that related deployments use the same terminal page after beginning with an investment-themed lure. Because the investigation did not observe any follow-on interaction, however, the purpose of the page and any subsequent activity remain unconfirmed.

© 2026 Abnormal AI Inc. All rights reserved.

A Live Data Relay in a Human-Controlled Attack Flow

ZeroTokens does not proxy the target's session to the real financial institution. The kit contains no genuine bank hostnames or references to browser automation tooling such as Puppeteer, Selenium, Playwright, headless browsers, or WebDriver. The platform's `form_command.navigate_to` functionality is wired to ZeroTokens' own relay infrastructure rather than to a financial institution, ruling out an automated AiTM proxy or session token relay within the kit itself.

Because ZeroTokens does not proxy the genuine site, it cannot automatically reproduce whatever authentication or verification step the institution presents. Instead, it relies on preconfigured institution-specific workflows that the operator can steer in real time.

The Live Data Channel

To coordinate the target-facing interaction in real time, the phishing page maintains a persistent WebSocket connection with the ZeroTokens platform for the duration of the session. On load, it connects to `wss://api.0tokens[.]xyz/ws/user/{project}/{session-id}`, while the operator console maintains a separate connection through the `/ws/admin/{project}/…` path.

The ZeroTokens backend relays information between those two connections. The phishing page reports session state to the platform, including whether the session is active, the target's current stage, and field values as they are entered. Commands issued through the operator console—including `navigate_to`, `input_message`, and `clear_row_state`—are relayed back to the target-facing page. Nothing is screen-shared or remotely controlled; the phishing page is ZeroTokens' own code and reports its state directly to the platform.

This telemetry is surfaced in the live session board, where each active target appears as a separate row, with columns corresponding to the information that particular phishing project is designed to collect (Fig. 14). Those fields populate as the target enters data; the project's harvest schema also tracks the target's current page.

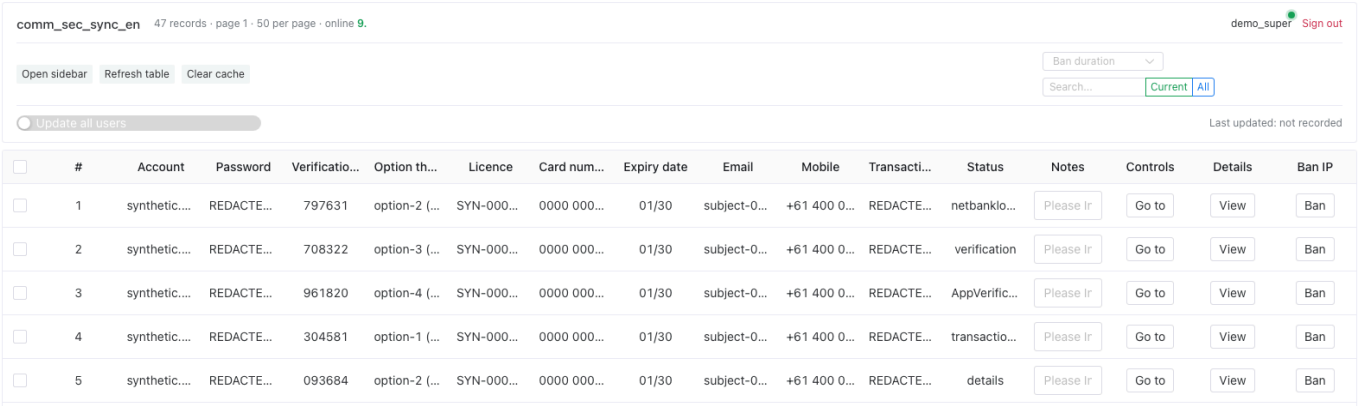

Fig. 14: The session board shows information entered into each live phishing session and the target's current page.

© 2026 Abnormal AI Inc. All rights reserved.

Real-Time Session Control

ZeroTokens' live control turns a set of preconfigured phishing screens into an adaptable workflow. The operator does not need to commit to one sequence in advance; they can respond to conditions within an active session—for example, repeating a failed step, switching verification methods, or introducing a custom prompt—while keeping the interaction within the same phishing flow (Fig. 15).

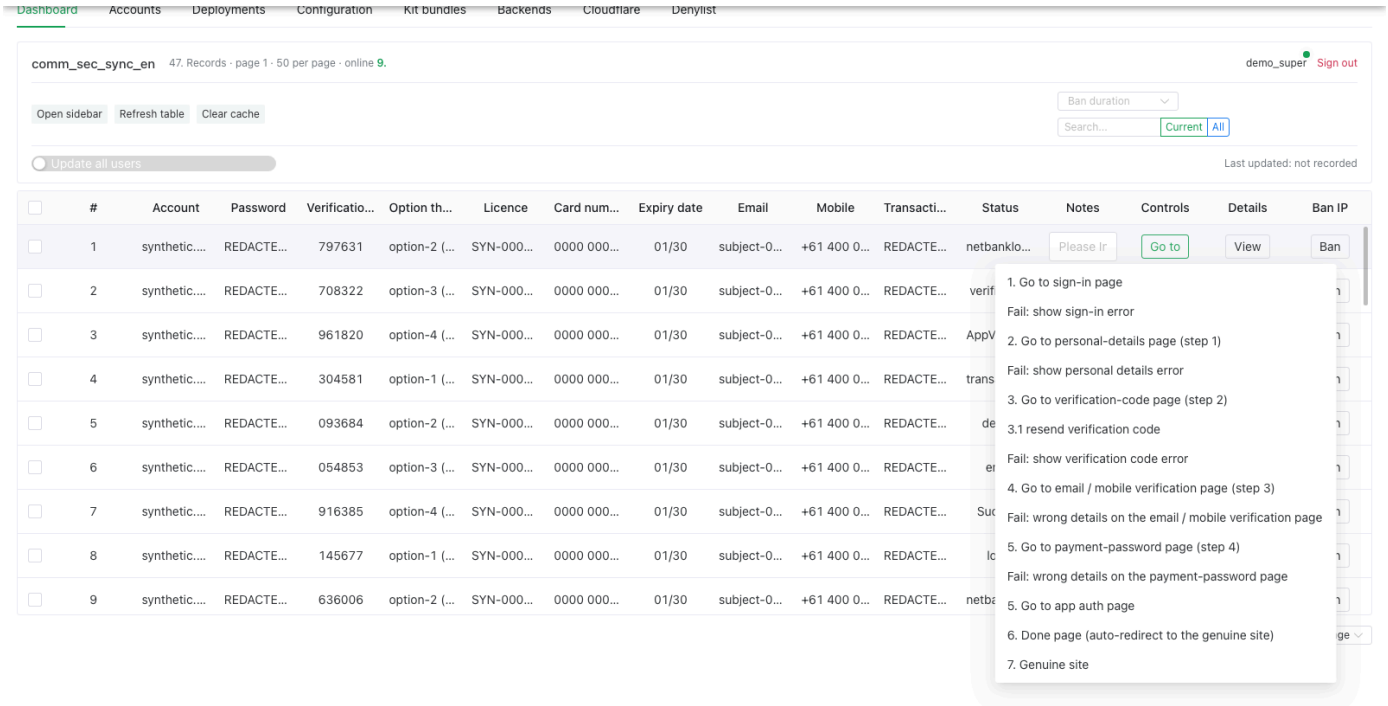

Fig. 15: The per-session command menu controls which stage or response the target sees next.

If a verification step fails, the operator can issue a corresponding failure command that displays an error, keeps the interaction open, and prompts the target to try again. This gives the operator another opportunity to capture a usable verification code or response—a capability most commodity phishing kits lack.

ZeroTokens can also switch from mobile verification to email when a mobile code cannot be obtained or fails. In the CommSec project, the operator can fall back to a dedicated email/mobile verification stage, with its own failure response. Similar email-code handling appears in other deployments, indicating that this is a broader platform capability, not a CommSec-specific feature.

The operator is not limited to the predefined stages of the phishing flow. Some session-board columns are configured as operator inputs rather than stored fields; whatever the operator enters appears as a prompt on the target-facing page. In one German template, an input can also accept a pasted image, enabling the operator to push a scannable photoTAN code into the live page. Together, these controls let operators introduce custom prompts dynamically as the session unfolds.

Once the collection process is complete, the operator can redirect the target to the institution's actual website. Ending the interaction on a legitimate site can make the transition away from the phishing flow appear less suspicious.

© 2026 Abnormal AI Inc. All rights reserved.

Using Harvested Data in Real Time

What happens after ZeroTokens collects credentials and verification responses is less certain. The platform contains no automated connection to the genuine institution, and whether an operator manually uses the harvested information there remains unconfirmed.

The architecture nevertheless can allow the operator to coordinate the phishing session with activity at the genuine institution in real time. Once a target's username and password appear in the session board, for example, an operator can use them to sign in to the real bank. If that login triggers an SMS verification code, the operator can direct ZeroTokens to display the corresponding OTP prompt. When the target enters the code on the phishing page, it appears in the session board, allowing the operator to enter it into the real bank session before it expires.

MFA Implications

ZeroTokens is built around authentication methods that rely on information or actions a target can provide while the phishing session remains active. The operator can request one-time passcodes—including SMS and email codes—or app-push approvals while they are still valid, allowing the platform to defeat MFA methods that depend on something the target can read, re-enter, or tap to approve.

That capability does not extend to phishing-resistant authentication. FIDO2/WebAuthn security keys and platform passkeys bind the authentication challenge to the legitimate origin, leaving nothing the target can read, copy, or relay.

© 2026 Abnormal AI Inc. All rights reserved.

ZeroTokens Scale and Operations

Behind the target-facing phishing flow, ZeroTokens provides a centralized control plane to configure institution-specific projects, deploy reusable phishing sites, manage infrastructure, and supervise live sessions. The captured console shows how those components are organized into a repeatable system spanning dozens of financial brands and multiple markets. Campaign-wide analysis subsequently confirmed phishing activity impersonating 25 financial brands, demonstrating that the platform was deployed across multiple brands and regions rather than merely configured to support them.

Institution-Specific Project Configuration

The console build captured on August 17 contained 53 financial institution projects across 14 locales, including banks and brokerages in Australia, Japan, South Korea, Germany, India, Canada, the United States, and other markets. Fifty-two of those projects were marked as live-relay enabled. Following the September 18 migration from `0tokens[.]xyz` to `bfovb[.]xyz`, the rebuilt console contained 58 brand projects, including five that were absent from the August 17 capture: Nabtrade, BMO, Consorsbank, Fineco, and ACOM.

Console Grouping	Brands Represented
Australia	CMC Markets Stockbroking; CommBank Share Trading; Nabtrade*; SelfWealth; Stake; TD; Totality
Canada	BMO*; National Bank Direct Brokerage; Qtrade; Questrade; RBC Direct Investing; Wealthsimple
Switzerland	Swissquote
Germany	comdirect; Consorsbank*; ING
English/Global	Binance; Capital.com; CommSec; eToro; IG; Interactive Brokers; moomoo; Robinhood; Saxo; SoFi; Tiger Brokers; Trade Republic; Webull
Europe	EXANTE;
India	Angel One; INDmoney; LCL Direct
Italy	Fineco*
Japan	ACOM*; Monex; Rakuten Securities; SBI Securities; SMBC Mobit; SMBC Nikko
South Korea	Kiwoom Securities
Malaysia	CGS International; M+Global
Singapore	FSMOne
Thailand	Bualuang Securities
United Kingdom	AJ Bell; Hargreaves Lansdown; Interactive Investor; Trading 212

© 2026 Abnormal AI Inc. All rights reserved.

United States	Ally Bank; Avenue Securities; E*TRADE; Fidelity; Merrill Lynch; Public.com; TradeZero
Unscoped	PayPal Sync

** Project added after the August 17 console capture and observed in the rebuilt September 18 console.*

Each project has its own harvest schema and operator command menu tailored to the institution's verification process. The harvest schema determines which fields appear on the session board, while the command menu determines which pages the operator can push to the target.

The differences between projects show how specifically those workflows can be modeled. For example, `kiwoom_sec_sync_kr` includes cloud-certificate fields, a 12-digit authorization code, and separate SMS-receive and SMS-reply fields. `comdirect_sync_de` includes photoTAN and mobileTAN, while `rakuten_sec_sync_jp` includes an emoji-grid selector. ZeroTokens can therefore reproduce institution-specific requests and verification steps across multiple markets instead of relying on a single generic credential form. The CommSec project analyzed in this report is itself one of the platform's more extensive configurations, with 13 harvest columns and 14 operator commands; only the Kiwoom Securities project has a wider harvest schema in the captured inventory.

The platform separately supports 36 card-issuer brand keys—35 Japanese issuers and AmEx US—through a simpler collection model. These brands have no individual project configuration and instead resolve to the shared `Default` template, which shows only a card number alongside a free-text note field for the operator.

The platform's broader card field vocabulary also includes fields for cardholder identity, expiry, 3D Secure verification, phone numbers, and address information, allowing more extensive cardholder data collection where a flow calls for it. This gives ZeroTokens two distinct collection models within the same control plane: institution-specific workflows for banks and brokerages, and a generic card-harvesting template reused across multiple issuers.

Templated Site Deployment

ZeroTokens centralizes management of its phishing sites through the same console used to oversee the broader operation (Fig. 16). The deployment panel shows active phishing sites, their infrastructure and configuration status, and the number of live sessions associated with each deployment.

© 2026 Abnormal AI Inc. All rights reserved.

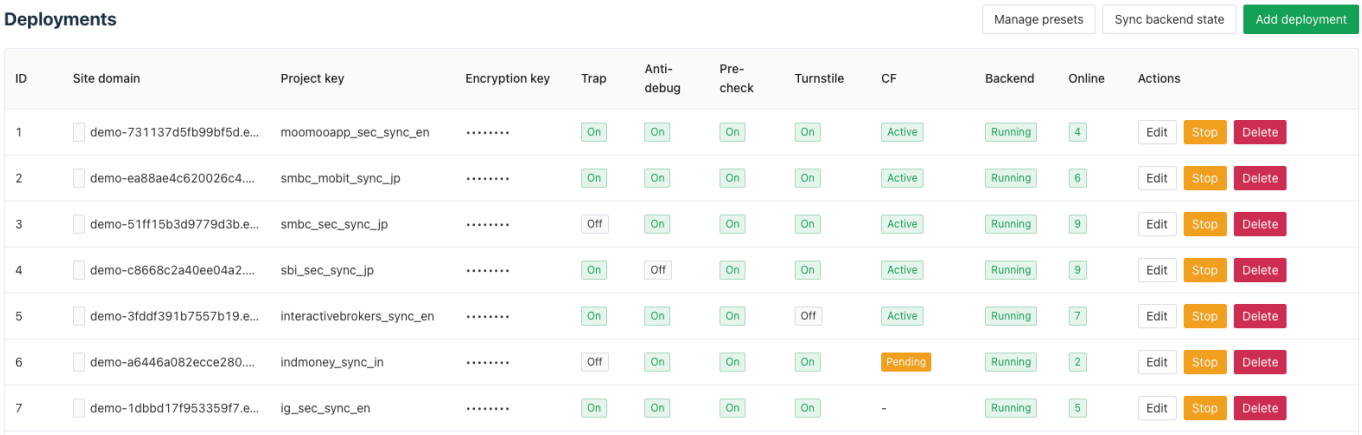

Fig. 16: The Deployments panel shows active phishing sites, their configuration and infrastructure status, and live session counts.

The platform also maintains a reusable library of phishing-site source code (Fig. 17). Operators can upload site packages to the console as ZIP archives and retain them for future deployments, reusing an existing phishing site instead of building each deployment from scratch.

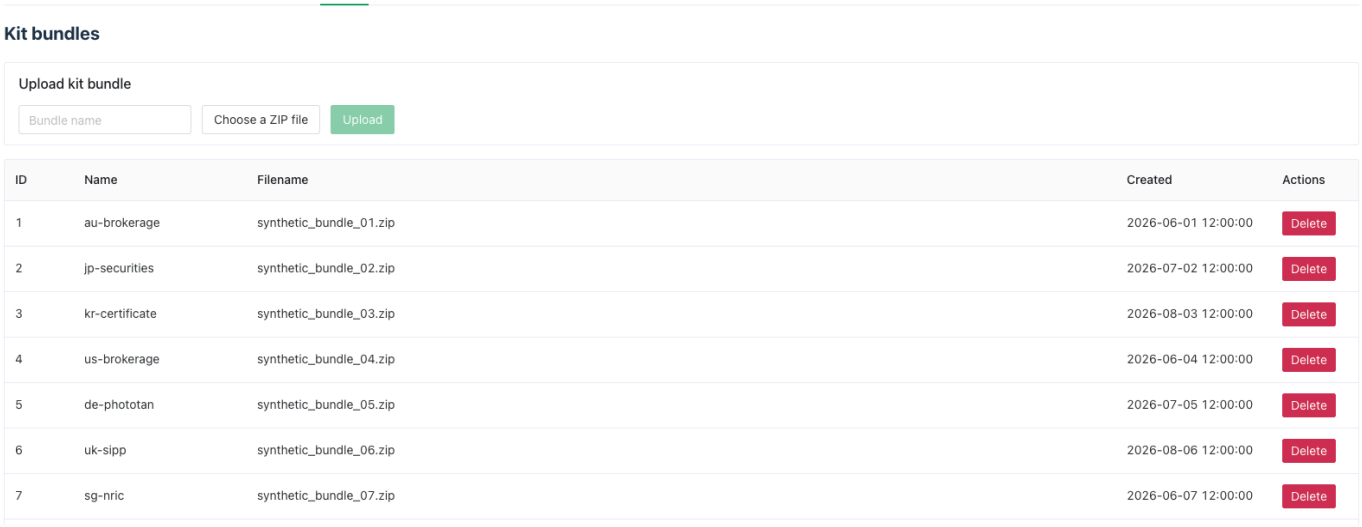

Fig. 17: The kit library stores reusable phishing-site source packages. Archive names shown are synthetic placeholders, not operator filenames.

When creating a deployment, the operator can configure the major components from a single form, including which stored site package to deploy and which institution-specific project to run (Fig. 18). The form also includes the site domain, backend, CDN account and country gating, Turnstile keypair, relay endpoint and `encrypted_key` value used for relay wiring, and four evasion controls: trap mode, anti-debug, pre-check, and Turnstile. Operators can also configure relay timing values, with presets available for reuse.

© 2026 Abnormal AI Inc. All rights reserved.

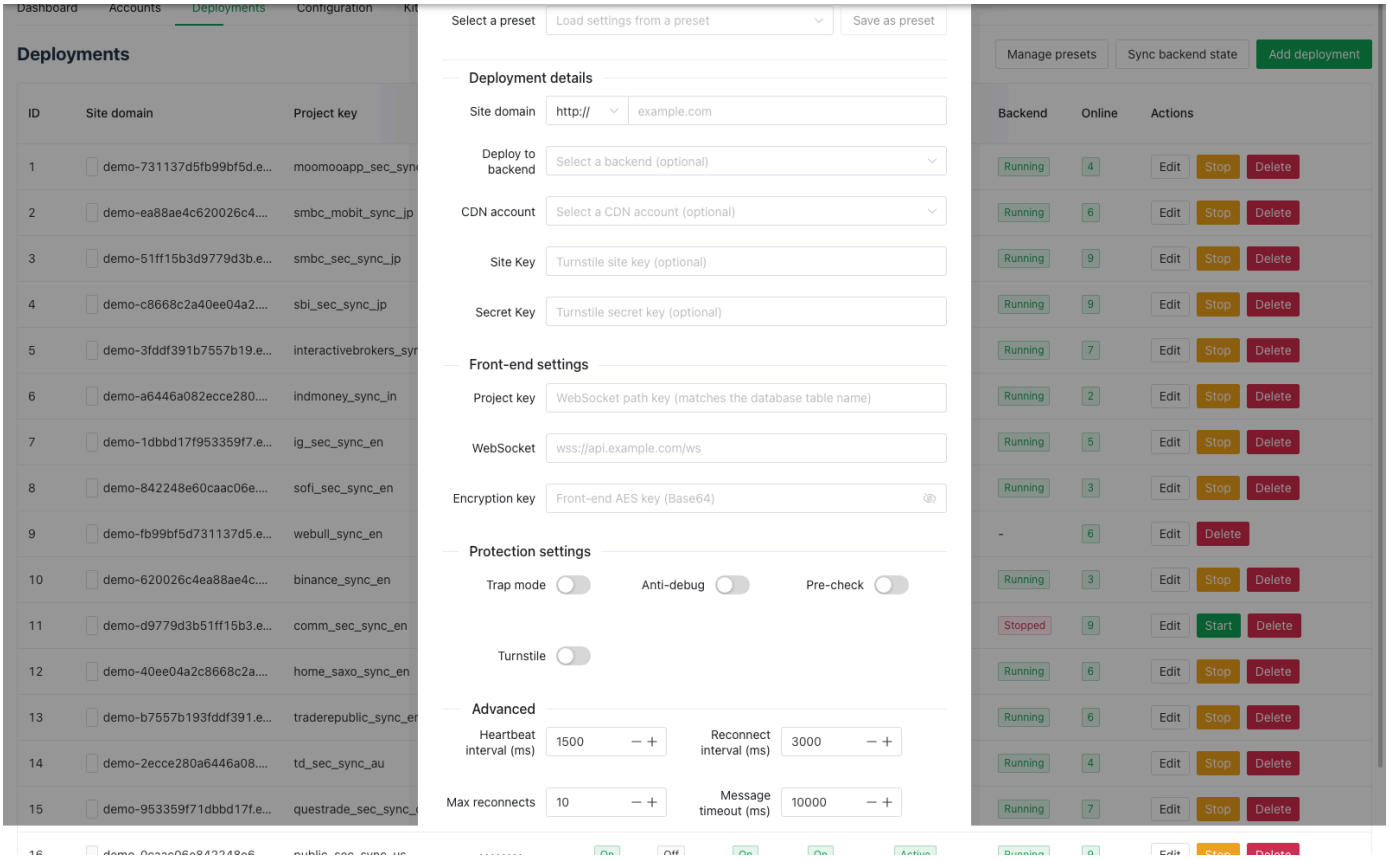

Fig. 18: The deployment form combines site infrastructure, collection settings, relay configuration, and evasion controls in a single interface.

Because those components can be drawn from existing site archives, backends, and CDN accounts, standing up a new phishing site becomes a largely templated process. Individual lure hosts are therefore relatively easy to replace, and taking down a single host does little to disrupt the broader operation.

Account Structure and Operator Assessment

ZeroTokens uses a simple account model with only two roles: `super_admin`, which has access to the broader control plane, and `operator`, which is limited to the session board (Fig. 19). Researchers found no billing, account-expiry, licensing, or tenant-separation functionality anywhere in the console. The console's own control channel uses a single static Advanced Encryption Standard (AES) key baked into the build, and exposes no marketplace, vendor handle, or support channel.

© 2026 Abnormal AI Inc. All rights reserved.

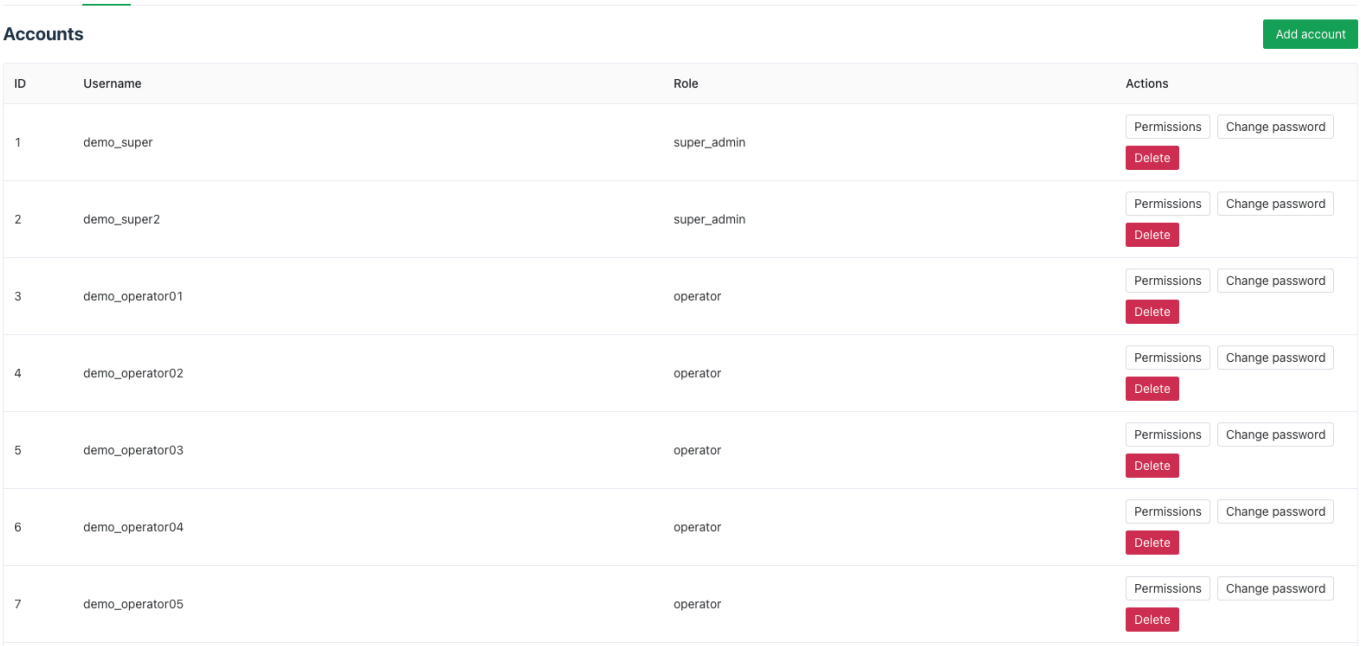

Fig. 19: The account schema contains only an ID, username, and role, with separate `super_admin` and `operator` access levels.

Those characteristics more closely resemble an internal staffing model than a customer-facing service, supporting a moderate-confidence assessment that ZeroTokens is in-house tooling rather than a rented platform. That assessment has an important limitation, however: from inside the platform, a tool operated privately by one group would be difficult to distinguish from one shared among a small set of trusted operators. The console therefore supports an assessment about the platform's operating model, but not the identity or precise structure of the group behind it.

No operator identity, persona, handle, messaging channel, wallet, or payment path appears anywhere in the control plane. The interface contains 863 Chinese-language UI strings, including an unreplaced placeholder page title, but language of authorship is the only origin-adjacent evidence identified. Because the platform targets institutions across Japan, South Korea, Germany, the United Kingdom, the United States, Australia, and other markets, the target set does not support geographic attribution. Chinese-language authorship alone is therefore insufficient to identify the operator’s location or a specific actor.

Built for Collection, Not Cash-Out

ZeroTokens is built to collect information rather than act inside a target’s financial account. Researchers reviewed the console’s complete action surface of 46 endpoints, all of which either manage the phishing operation—including accounts, sites, CDN accounts, backends, configuration, and denylisting—or provide access to harvested data. None initiates activity in a victim’s account. A separate sweep of all 54 template definitions in the August 17 captured console build—53 institution-specific projects plus the shared `Default` template—403 field definitions, and 382 operator commands likewise found no functionality for withdrawals, transfers, adding payees or beneficiaries, checking balances, or placing buy, sell, or order actions.

© 2026 Abnormal AI Inc. All rights reserved.

Additional evidence reinforces that distinction. The only money-related string identified was “payment password,” which refers to a credential, not a transaction. The templates contain no balance, holdings, buying-power, or instrument fields, and every terminal command either marks the interaction as Done or directs the target to the genuine institution. Although searches for `transfer` and `order` returned matches in the console, those occurrences were CSS terms with no financial function. Those findings indicate that ZeroTokens has no cash-out capability.

Any subsequent account activity would therefore occur outside the platform, through the genuine financial institution. The ultimate objective remains unresolved. The available evidence supports several plausible outcomes, including session takeover combined with identity harvesting, brokerage account abuse, and theft followed by liquidation. The SpaceX IPO exit page also suggests that some flows may be intended to set up a separate follow-on investment fraud approach. Because none of these downstream outcomes were observed directly, the platform’s harvested field set and terminal views indicate possible monetization paths rather than confirmed financial activity.

© 2026 Abnormal AI Inc. All rights reserved.

Campaign Scope and Evolution

ZeroTokens first surfaced through an Australian wave impersonating CommSec, but a full remeasurement across Abnormal message telemetry shows that activity was one part of a much broader operation. From April 21 through September 25, 2026, Abnormal observed at least 1,222,567 messages linked to the operation across 2,814 customers, impersonating 25 brokerage and banking brands. Canada and Australia accounted for most observed message volume, with smaller populations across other markets.

That total should be treated as a conservative floor. The operation used multiple delivery and lure-hosting patterns over the five-month period, including disposable random `.com` domains that cannot be enumerated exhaustively without introducing unrelated legitimate mail. Of that total, 652,452 messages belong to a backend-confirmed subset on `.live` lure hosts tied to the operator backend. The SendGrid and random `.com` activity was attributed through converging account, sender-pattern, brand, and registrar evidence. The broader campaign therefore extends well beyond the August activity that initially surfaced the ZeroTokens platform.

Campaign Generations and Scale

The operation evolved through three overlapping lure infrastructure generations rather than a single fixed delivery model.

Generation	Window	Delivery Model	Observed Volume	Lead Brands
Gen 1	Apr 21–early Jul	Long random `.live` lure hosts associated with an earlier backend	~369,000	CMC Markets, CommSec, TD, National Bank, Westpac, SelfWealth, Totality, FundSupermart, Avenue
Gen 1.5	Jul 20–Aug 10	Random `.com` lure hosts	~341,000	Wealthsimple, Interactive Brokers, TD
Gen 2	Jul 26–Sep 25	18-character `.live` lure fleet and abused SendGrid accounts	~512,000	moomoo, Interactive Brokers, CommSec, National Bank, Nabtrade, BMO, Interactive Investor, Saxo, Swissquote, Consorsbank, Fineco

Observed volume increased from approximately 28,000 messages in April to 109,000 in May, 232,000 in June, 236,000 in July, and 405,000 in August, with approximately 212,000 additional messages observed during the partial September measurement period. The operation entered a roughly two-week lull from June 29 through July 12, which coincided with no new operator domain registrations and the transition to the infrastructure later associated with `0tokens[.]xyz`. This is consistent with an infrastructure change rather than the end of the operation.

© 2026 Abnormal AI Inc. All rights reserved.

W-8BEN tax recertification was the operation's most commonly observed pretext, accounting for approximately 478,000 messages, or about 39% of campaign volume. Other activity used account security, device enrollment, two-factor authentication activation, margin call, and pre-IPO themes, demonstrating that the operation was not dependent on a single social engineering narrative. Across a separate 182-day hunt for W-8BEN-themed mail, 92% of identified messages were classified as attacks, and at the August peak, ZeroTokens-linked activity accounted for roughly seven in ten attack-marked W-8BEN messages observed.

Victimology and Targeting

Across the full campaign, Canada and Australia accounted for at least 88% of observed message volume. Canada represented 47.9% and Australia 40.1%, while 7.7% of messages went to generic `.com` recipient domains whose geographic location could not be determined from the address alone. Smaller identified populations were associated with the United Kingdom, Singapore, Switzerland, and other European and Asian markets. Because geographic attribution is based on recipient domain information, the Canada and Australia percentages should be treated as minimums rather than precise measures of the underlying target population.

The campaign affected organizations across a broad range of industries. Education represented the largest number of affected organizations, followed by manufacturing, business services, finance, retail, construction, government, insurance, software, and healthcare. Message volume was more concentrated than organization count, however. Rather than selecting individual targets, the operator appears to have distributed lures broadly across organizational mail directories, allowing a handful of large environments to account for disproportionate message volumes. One Canadian university alone received approximately 463,000 messages, or 37.8% of observed campaign volume. Organization count is therefore a better measure of the campaign's sector breadth than raw message volume.

Targeting was similarly broad by employee seniority. Among recipients whose job titles were known, representing about 73% of observed mail, approximately 62% were individual contributors and 30% were managers or team leads. Directors represented roughly 5%, while C-Suite, vice president, and board-level recipients together accounted for only about 2.9%. That distribution indicates that the campaign was not narrowly focused on executives; the operation cast broadly across employee seniority rather than concentrating on leadership roles.

Australian CommSec Wave Victimology

The Australian CommSec wave shows how that broader targeting appeared within an individual deployment. Organizations associated with `.au` domains accounted for 604 of the 732 affected organizations and received 42,398 of the 45,774 messages—82.5% of affected organizations but 92.6% of total message volume. Insurance represented the largest share of affected organizations, while education received the highest message volume, with 11,927 messages across 85 organizations. Energy and utilities received 4,374 messages across 20 organizations, and minerals and mining received 3,193 across 17, resulting in comparatively high volumes relative to the number of organizations affected.

© 2026 Abnormal AI Inc. All rights reserved.

Earlier Related Activity

Public reporting on related W-8BEN phishing predates the August investigation and provides evidence of continuity in the operation's associated sender infrastructure. Interactive Brokers customers were targeted with the same W-8BEN pretext in October 2025, while similar activity targeting CommSec customers was publicly documented in June 2026.

One sender domain used in the October 2025 activity, `xbvdsi[.]com`, later appeared in activity associated with the broader ZeroTokens operation. Registration data shows that the domain was registered on August 15, 2025, never lapsed, and used the same registrar as the ten main-wave sender apexes. That reuse establishes continuity in associated sender infrastructure dating back to at least October 2025. It does not establish that the ZeroTokens live control platform itself was in use at that time. Earlier public reporting describes simpler credential-focused phishing and does not document the live, operator-controlled workflow identified through the August 2026 platform capture.

Fleet Evolution and Platform Migration

A September 2 reevaluation showed how quickly individual ZeroTokens lure hosts could disappear even while more durable parts of the operation remained active. Researchers enumerated 47 lure apexes through the control plane pivot, 34 of which were also identifiable through a shared `config.js` body hash. By that date, only five of the 47 lure apexes still resolved, and just two continued to serve ZeroTokens content: a Nabtrade deployment delivered through Cloudflare and a Saxo deployment exposed directly from the hosting origin.

The reevaluation also showed that DNS resolution alone was an unreliable measure of whether a lure remained active. Three of the five resolving hosts no longer served the kit, while virtual hosts that had been removed from the shared origin instead produced TLS errors. The origin itself remained active and continued serving configured fleet hosts based on Server Name Indication.

The gate had also changed. Whereas the captured August build relied on a client-side verification flag, surviving deployments returned an HTTP 400 response when a Turnstile token was absent, indicating that at least some deployments had moved Turnstile validation server-side.

A later infrastructure event reinforced the distinction between replaceable lure hosts and the broader platform behind them. On September 18, `0tokens[.]xyz` went under registry hold. Within roughly 20 hours, the operator had rebuilt the console on `bfovb[.]xyz`. The rebuilt environment also contained five brand projects absent from the August 17 console capture: Nabtrade, BMO, Consorsbank, Fineco, and ACOM. Nabtrade had already been observed independently during the September 2 reevaluation, while the other four provide additional evidence that the platform continued to evolve after the original capture.

© 2026 Abnormal AI Inc. All rights reserved.

Conclusion

ZeroTokens demonstrates how a phishing operation can become more capable without relying on a novel exploit. Its defining feature is the operational system built around the phishing page: a human operator watching each target's session, reviewing submitted information, and deciding which institution-specific prompt appears next. That turns credential harvesting into a live, adaptive interaction rather than a fixed sequence of forms.

The platform's scale shows this model is not limited to a single brand or campaign. By September 2026, ZeroTokens supported 58 financial brand projects, while campaign-wide analysis identified phishing activity impersonating 25 brands across more than 1.2 million messages. The operation also proved resilient, rotating sender and lure infrastructure and rebuilding its control plane on new infrastructure within roughly 20 hours after the original backend was placed under registry hold.

ZeroTokens is not a traditional adversary-in-the-middle platform. The captured code contains no automated connection to a genuine financial institution and no mechanism for stealing authenticated session tokens. Instead, it is designed to collect credentials, identity information, one-time codes, app approvals, and institution-specific authorization secrets while the target remains engaged. Whether operators manually use those inputs against the genuine institution remains unconfirmed.

For defenders, the broader lesson is to focus on the behaviors that persist when infrastructure changes. Phishing-resistant authentication can close the relay path for user-transferable MFA factors; detection and response should prioritize recurring platform signatures, operator-controlled infrastructure, and the full scope of data exposed during a session—not credentials alone.

© 2026 Abnormal AI Inc. All rights reserved.

Detection, Hunting, and Response

ZeroTokens makes many of its most visible components relatively easy to replace. Sender domains and lure hosts can rotate, and even backend and control-plane infrastructure can be replaced. Site packages, project configurations, and other deployment components can also be reused or redeployed. Detection and disruption should therefore focus on recurring characteristics instead of any single phishing page, pretext, or hostname.

The platform also collects more than login credentials. A single interaction can expose authentication responses, identity information, and institution-specific authorization credentials, making it important to recognize the activity as a multi-stage harvesting session rather than conventional credential phishing.

Email and Delivery Signals

The Australian CommSec wave combined legitimate email infrastructure with attacker-controlled sender domains. Messages were delivered through abused SendGrid accounts, carried SendGrid DKIM signatures, and passed SPF, DKIM, and DMARC checks. The sender domains were aged for several months before use, reducing some of the obvious indicators associated with newly registered phishing infrastructure.

Several recurring behaviors provide stronger hunting leads:

- **Send-only domains with no MX records.** None of the ten sender domains used in the Australian CommSec wave published an MX record, indicating that they were configured to send mail but not receive it.
- **Newly active senders using a single subject template at volume.** Each of the ten sender domains in the Australian CommSec wave was associated with one subject template, with no rotation between domains. That one-to-one relationship can help connect a known sender to additional messages using the same template, or identify associated infrastructure after a template is discovered.
- **Persistent email service provider identifiers.** Provider-specific account identifiers can help connect activity even when the attacker changes the sending domain. In this operation, nine abused SendGrid customer IDs provide reusable pivots that can persist across sender-domain rotation.
- **Anomalous submission HELO strings.** The Australian CommSec wave used 12 forged HELO strings: three did not resolve, while the other nine belonged to unrelated third parties. Where that telemetry is available, those inconsistencies provide another useful hunting signal.
- **Recurring lure-fleet characteristics.** One major ZeroTokens delivery generation used 18-character random-alphanumeric `.live` domains registered through GNAME.COM, fronted by Cloudflare, and provisioned in closely timed batches shortly before use. The combination can help identify related infrastructure even as individual lure hosts rotate.
- **Multi-stage collection from the same client.** Requests for license details and a trading password as separate steps minutes apart from the same client can indicate an active, operator-controlled harvesting session rather than a one-time credential form.

© 2026 Abnormal AI Inc. All rights reserved.

Validate sender-domain pivots before promoting associated domains to operator-controlled indicators. In the broader campaign hunt, a raw telemetry pivot also surfaced legitimate third parties and victim organizations that had forwarded the phishing internally. Treat sender relationships as hunting leads until ownership and role are established.

The lure itself should carry less weight. W-8BEN recertification was the most common observed pretext in the broader ZeroTokens-linked campaign, but it is reusable and has appeared outside this operation. The pretext can support hunting for phishing activity, but it does not, by itself, isolate this operator or platform.

Platform-Level Signatures

ZeroTokens exposes several recurring backend and network behaviors that are more closely tied to how the platform operates. These include its encrypted request wrapper, persistent victim-side WebSocket connection, runtime configuration calls, source package upload workflow, and consistent HTTP response behavior.

Behavior	Role	Why It Matters
Request bodies use a single JSON member, `data`, containing `base64(IV[16] ‖ AES-256-CBC ciphertext)`; the bearer token is encrypted under the same static key; responses are cleartext	Console ↔ backend	The same key is used across deployments, making this request shape a platform-level network signature rather than a campaign-specific artifact
Separate victim-side (`/ws/user/...`) and operator-side (`/ws/admin/...`) WebSocket connections; the backend relays target input and operator commands between them	Live relay	The backend relays victim input to the operator and operator commands back to the phishing page, enabling real-time control of each active session.
`GET /turnstile/site-key/` returns the Turnstile sitekey at runtime	Central gate configuration	The sitekey is delivered dynamically and is not embedded in the page, making the request itself a useful platform signal
Multipart `POST /source/upload/` accepts a kit ZIP that is referenced by `source_code_id` when a deployment is created	Kit distribution into the deployment library	Phishing-site source is maintained as reusable inventory that can be redeployed across lure hosts
`POST /site/frontend-config/` returns per-deployment values including `projectName`, `wsUrl`, `encryptedKey`, `heartbeatInterval`, and the four evasion toggles	Victim-side runtime configuration	Core relay and evasion settings are delivered as deployment configuration instead of being fixed entirely in the phishing page

© 2026 Abnormal AI Inc. All rights reserved.

Requests to an unmatched path return the same 467-byte SPA body (`cd7f93ff...`), observed across 84 probe paths	HTTP fingerprint	The consistent response provides a content-level fingerprint that does not depend on a specific lure URL
Requests using an unconfigured Host return HTTP 200 with a counterfeit 130-byte nginx 404 body (`7a50dcaa...`)	HTTP fingerprint	The response only appears to be a 404; the body and HTTP status provide a repeatable way to identify the deployment.
Configured HTTP lure hosts return a consistent 166-byte redirect body (`87a9323a...`) when redirecting to HTTPS	HTTP fingerprint	Complements the SPA and counterfeit-404 body hashes when classifying ZeroTokens fleet infrastructure

These signatures describe the platform itself, not merely the infrastructure placed in front of a target. Individual lure domains can disappear quickly, while the live-relay mechanics, runtime configuration requests, request wrapper, and deployment workflow reflect recurring properties of ZeroTokens.

Disruption Priorities

ZeroTokens makes individual phishing sites relatively easy to replace. Operators can combine previously uploaded site packages with existing backends, CDN accounts, project configurations, and relay settings to create new deployments without rebuilding the infrastructure stack from scratch. Taking down one lure host can therefore remove an active site without materially disrupting the underlying platform.

Where possible, defenders should prioritize infrastructure and identifiers that sit higher in that stack: confirmed ZeroTokens backend and control-plane infrastructure, abused email service provider accounts and account identifiers, recurring platform signatures, and other operator-controlled infrastructure. Do not block shared hosting origins, CDN ranges, or email-provider infrastructure wholesale, as this would affect unrelated legitimate services.

IOCs

This list is limited to operator-controlled infrastructure and high-confidence ZeroTokens platform fingerprints. Shared hosting, CDN, and email service provider infrastructure, as well as impersonated brands and recipient organizations, are excluded to avoid collateral damage from blocking.

Indicator	Type/Role	Action

`0tokens[.]xyz`
`api.0tokens[.]xyz`
`admin.0tokens[.]xyz`
	Platform apex; backend/live relay/gate-key service; operator console	Block

© 2026 Abnormal AI Inc. All rights reserved.

fobwj[.]com qtgep[.]com yigau[.]com zhengkanginternational[.]asia hojietrading[.]asia macautextile[.]asia kwanlunenterprise[.]asia afrasianfood[.]asia hopewellknittingfactory[.]asia chengdainternational[.]asia	Operator sender apexes	Block
ywqma[.]com yzraj[.]com ozkui[.]com mcpaj[.]com cjhwl[.]com kuoduoduo[.]com pciqz[.]com jwvnh[.]com	Canadian-wave sender domains	Block
dkprlh22dvbgmx1hfd[.]live d5e3d7stv2t9c6wghp[.]live fsqedvfdz2vviehpjx[.]live 7dbud7twop8kyjsx0p[.]live z1d5doj2xjq8q33ct8[.]live soj54718oahk25lli4[.]live w8l7smx4cfidedzeb2[.]live citg4xi2xcup5cpu84[.]live djsjb2tcu77lkzx9ay[.]live ecrml50ij3ccz0roi1[.]live t704dmq753cjd1jv0f[.]live	Lure fleet (captured or infrastructure-linked)	Block
46hwg7vtakoqm0daje[.]live p6t1w8d4j9f2h7s5ky[.]live a3v9n2m7x5q1k4r8dc[.]live z8f2p6w1j9t4g7b3mx[.]live 8390x8ms6vqr2bon5p[.]live z2f6w9p3j1t7g4b8mx[.]live ei787pr3hd117g1yfv[.]live xontmzggrk96w6s6a5[.]live d6p8w1j4t9f2h7s5ky[.]live	Canadian-wave lure hosts (shared infrastructure attribution)	Block
hxxps://api[.]0tokens[.]xyz/turnstile/site-key/	Runtime gate-key endpoint	Detect
N8BN4W@fobwj[.]com NO8MQKX1@zhengkanginternational[.]asia	Observed envelope senders	Block

© 2026 Abnormal AI Inc. All rights reserved.

`0x4AAAAAAEEW0ZNv6P07pBM9`	Operator Turnstile sitekey; platform fingerprint	Detect
`cd7f93ffa6ab1444…` — 467 B SPA fallback `7a50dcaa62ecb3ba…` — 130 B counterfeit nginx 404 `87a9323a…` — 166 B redirect	Fleet response-body hashes	Detect
阿巴阿巴	Operator-console UI string	Detect

`53411154`
`53447722`
`53457520`
`53477852`
`53507665`
`53562556`
`53632199`
`53648237`
`53648238`
	Abused SendGrid customer IDs	Report to ESP

Do not block shared hosting origin `190.123.46.44`; the infrastructure is multi-tenant and hosts unrelated legitimate services. Legitimate domains belonging to impersonated financial institutions are likewise excluded because they are not operator-controlled infrastructure.

© 2026 Abnormal AI Inc. All rights reserved.