Solidarity Playbook

Resisting six weeks of sustained phishing attacks

International Civil Society Centre

CyberPeace Institute

solidarityaction.network

1. Overview

Transparency International is a global movement working to end the injustice of corruption. Through research, advocacy, and campaigning, they work to expose the systems and networks that enable corruption to thrive, demanding greater transparency, accountability, and integrity at all levels and across all sectors of society. The organisation's structure includes national chapters in more than 100 countries. Each national chapter has its own IT department, and these departments coordinate with the Secretariat in Berlin, Germany, for different types of support.

In 2019, Transparency International experienced a sustained and sophisticated phishing attack, which was detected through a sharp increase in failed attempts to log in to organisational email accounts. Over a six-week period, the five-person IT team at the Secretariat monitored and responded to this persistent effort to breach the organisation's systems.

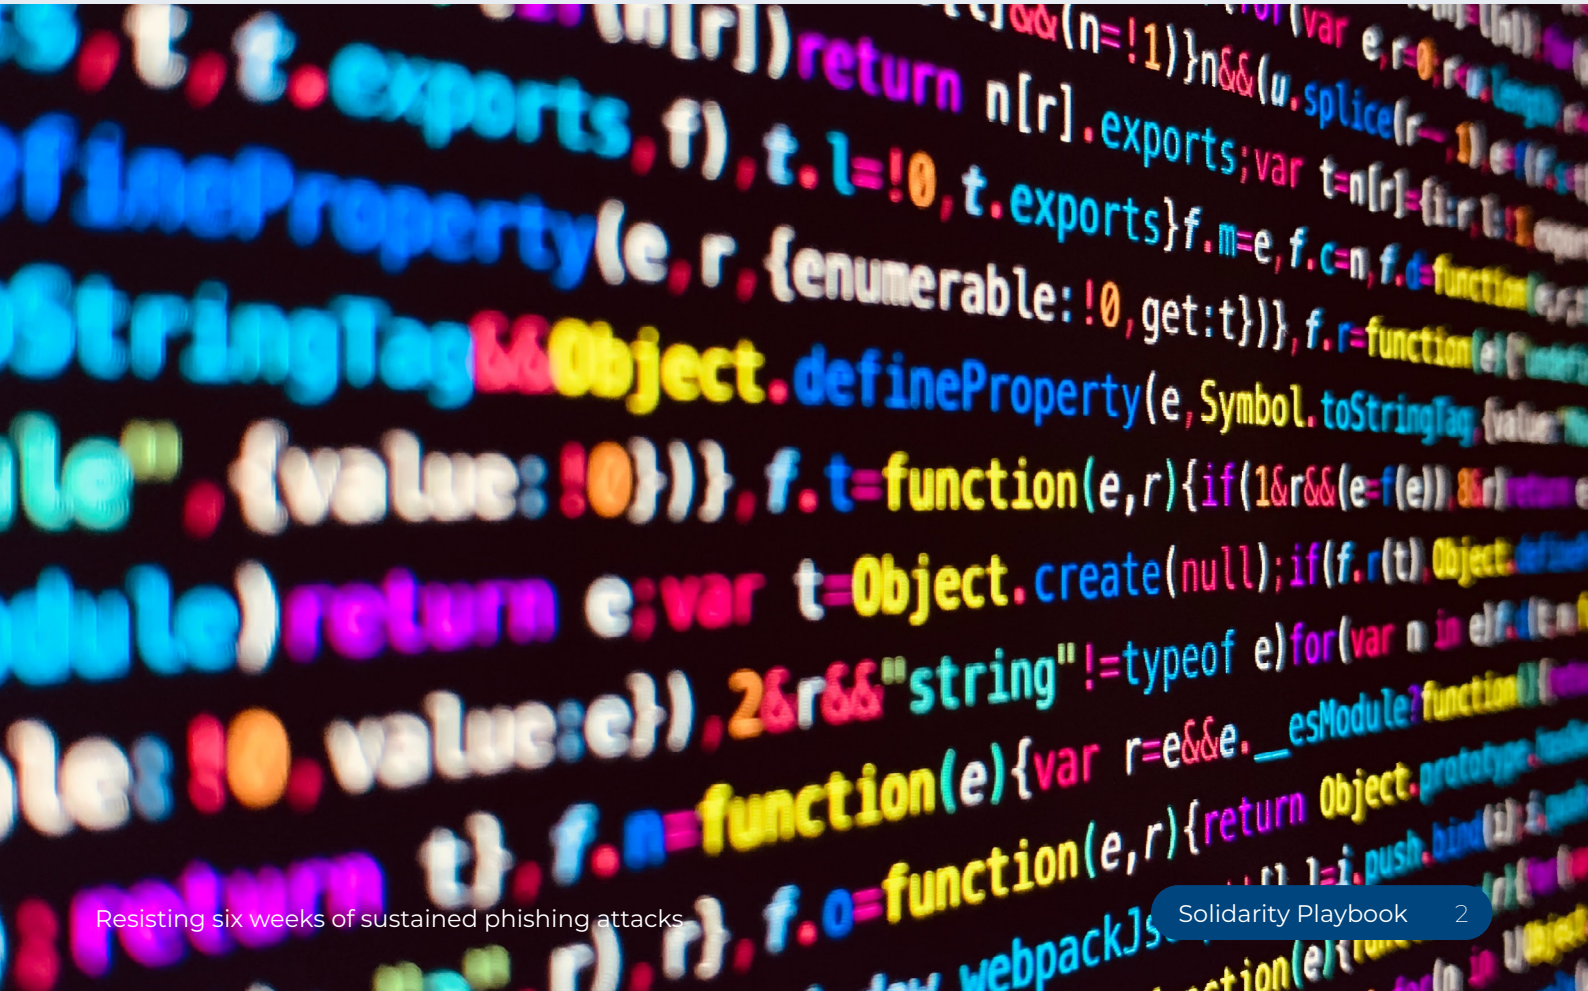

2. What happened?

In the summer of 2019, the IT team received an email from Microsoft alerting them to a **sudden and unusual increase in activity, resulting in failed attempts to sign in to the organisation's Microsoft Office accounts**. Transparency International uses Microsoft Office 365 services for email, browsers, and internal collaboration tools. While it was common for failed login attempts to occur (for example, due to someone in the organisation forgetting their password and accidentally using the wrong credentials or an external actor trying to enter the system through email or other services), the volume of failed attempts in this case was extraordinary and indicated a sustained effort to breach the organisation's IT system.

After receiving the email from Microsoft, the IT team did not know what to make of the warning. They were reassured that staff accounts were **protected by multi-factor authentication**, which felt like adequate protection from the barrage of login attempts, but they decided to continue to closely monitor these attempts. Whenever they saw an account that used the correct login credentials (such as the correct email address and password) but failed to apply the correct multi-factor authentication, the **IT team would warn the staff member whose account had been targeted and ask them to change their password**. At this early stage, the IT team decided against sending out a warning email to all staff to avoid creating fear or paralysing people's work.

MULTI-FACTOR AUTHENTICATION (MFA)

Multi-factor authentication is a layered approach to securing data and applications in which a system requires a user to present a combination of two or more credentials to verify the user's identity for login. MFA increases security because even if one credential becomes compromised, unauthorised users will be unable to meet the second authentication requirement and will not be able to access the targeted physical space, computing device, network, or database.

Cybersecurity and Infrastructure Security Agency (2022):
Multi-factor authentication (MFA)

One week after the initial warning email from Microsoft, senior management started to receive phishing emails impersonating the Chair of the Board and mimicking their writing style. The IT team realised that this was a **spear phishing attack** and, for the first time, understood that the organisation was being targeted. The phishing emails were sent from Gmail addresses but had a high level of sophistication, as they closely imitated the writing styles of previous emails that had been sent. For instance, on the sixteenth day of the cyberattack, the attackers sent an email mimicking a warning email that had been sent by the Head of Technology 24 hours earlier. In their email, the attackers offered additional support regarding the phishing emails by asking staff to click a link to get more

information about the ongoing attacks! Given how closely the phishing emails mimicked management emails, the IT team suspected that the attackers could see emails within the organisation's system – but they could not identify where or how the attackers were able to do so.

SPEAR PHISHING

A spear phishing attack is an advanced method of creating non-legitimate emails aimed at well-chosen targets within an organisation. Cybercriminals disguise themselves as trustworthy individuals and create stories which victims believe. Malicious actors typically prepare for an attack by getting to know the victim through social networks and gaining access to sensitive information via email spoofing or infiltrating other online messaging systems. Preparation can take weeks or months before an actual attack is carried out.

CyberPeace Institute (2020): ***Hackers Trick Humanitarian Non-profit into Big Wire Transfers***

The IT team continued to "firefight" the login attempts and phishing emails. In the third week, Transparency International received another email from Microsoft, which said that attackers had successfully breached the accounts of two staff members. This meant that the attackers had not only used the correct login credentials for these staff members but were able to "confirm" their identity by completing or bypassing the multi-factor authentication process (Transparency International used SMS-based multi-factor authentication at the time). The two staff members whose accounts were breached were not senior staff and fortunately did not have access to sensitive data. However, the organisation's IT security had been breached by cyberattackers. The IT team was able to end all active sessions and change passwords on the affected accounts to ensure the attackers were logged out.

After six weeks, the attacks subsided. To this day, Transparency International is not sure why the attacks stopped. It could be that the attackers found what they wanted, that it was too difficult to find and access what they wanted, or that there was a time limit set on the campaign against them.

Overall, the six weeks of attacks on their system were a demanding and strenuous time for the organisation. One of the two IT staff members directly responding to the incident had just joined the organisation three days prior to the start of the incident. IT staff did not know what was coming next and had to "firefight" this attack while still completing their regular work.

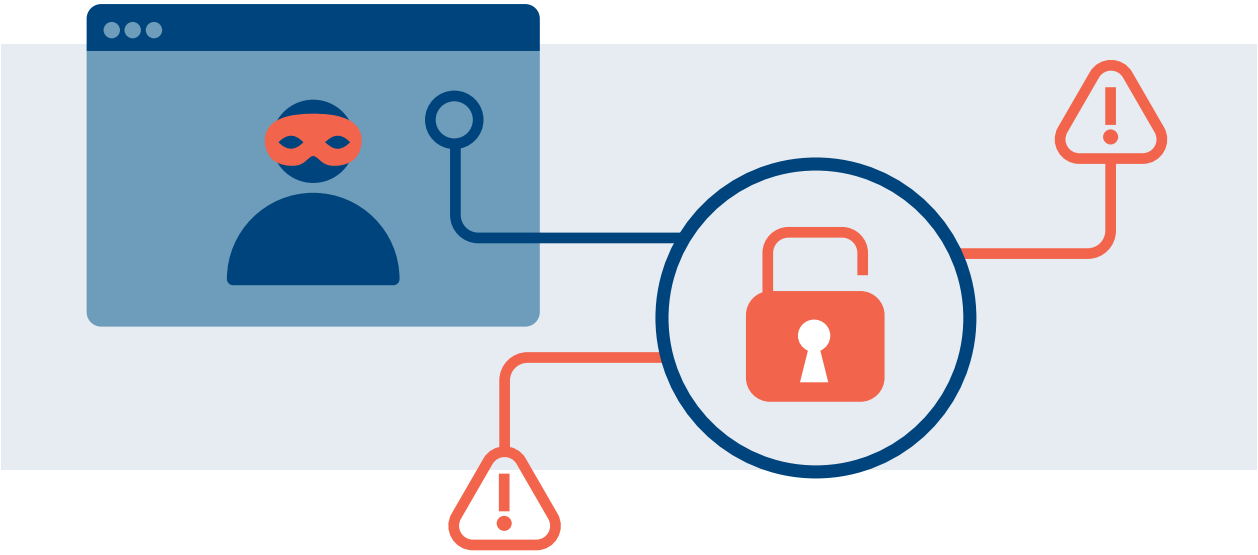

3. Response

During the first two weeks of the attack, the IT team was in constant communication with senior management but did not alert all staff. In meetings with senior management, they probed into why this attack was happening and what the attackers could possibly want. By the end of the second week, given the persistence of the phishing attacks, the head of IT sent out an awareness and warning email to all staff. The aim was to prevent staff from accidentally exposing their accounts. This was the email that the attackers mimicked in a phishing email sent to all staff on the 16th day of the incident.

Transparency International had a retainer relationship with a cybersecurity company that ran an intrusion detection service to constantly scan communications across the organisation's network for a breach. By day 13 of the attack, the company notified Transparency International that three devices in the Secretariat's network were communicating with external IP addresses. In their analysis, they noted that the IPs likely belonged to state-sponsored attackers who had substantial resources to put toward an attack. **Based on similarities to previous attacks, the cybersecurity company suspected (though they could not completely verify) that the attackers were a group known as Cozy Bear or APT 29**, a prolific group linked to Russian intelligence which "hunts for confidential information stored in the networks of governmental organisations, political groups, and think tanks."

The cybersecurity company also discovered that although the intrusion attempts seemed to originate from Chinese IP addresses, they were actually Russian IPs masked as Chinese addresses. They found that the attackers were **registering many domain names that were close to the domain names of several international civil society organisations (ICSOs)** – a type of cyberattack known as **domain spoofing**. Following the analysis by the cybersecurity company in which the organisation learned that this attack was from a state-sponsored actor, Transparency International became more serious in their response.

DOMAIN SPOOFING

A fake website name or email domain created to trick the user into sharing personal information (such as login credentials or credit card details) or downloading malware.

Solidarity Action Network (2022): ***Navigating cybersecurity: Guidance for (I)CSO professionals***

Cyberattackers often switch the position of letters (like vowels a and e) or exchange periods for dashes, which can easily be missed.

Some examples of domain spoofing that Transparency International has experienced are:

- transparancy.org instead of transparency.org
- mail-transparency.org instead of mail.transparency.org
- login-transparency.org instead of login.transparency.org

The IT team tried to stay ahead of the attackers by putting blocks on the firewall so staff would not accidentally open phishing links. "*We had a game of cat and mouse for a few weeks, as the attackers switched up [their approach] and we blocked*," said an IT staff member. Suspecting that their email communications were being monitored, the **IT team moved their communications offline and relied on in-person conversations to coordinate** to avoid further compromise. They also started blocking email addresses that were suspected of being related to the attackers and asked staff to be alert for phishing emails, providing a list of actions that staff should and should not take (such as to avoid clicking links or downloading files from unknown accounts and to report suspicious emails).

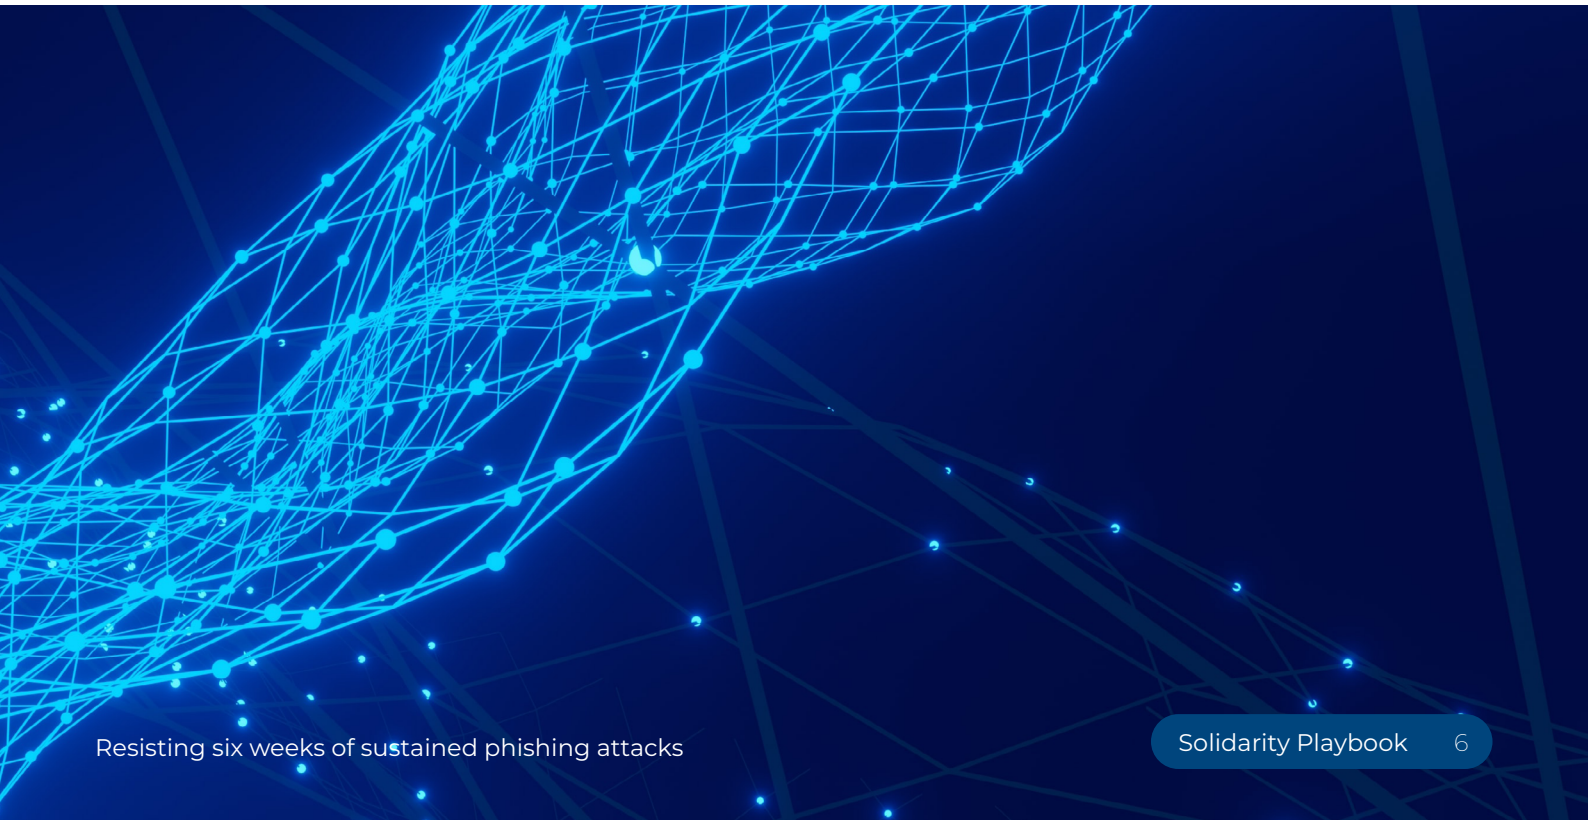

BEST PRACTICE FOR ALERTING STAFF TO PHISHING EMAILS

It is recommended to never share a sample phishing email by forwarding the email, as it could still contain malicious links or attachments. The best way to inform staff about phishing threats is to send a separate email with details and share a screenshot of the phishing email, which will be harmless.

For more advice, see Solidarity Action Network (2022):
Navigating cybersecurity: Guidance for (I)CSO professionals

Eventually, the attackers moved to European servers, probably to avoid detection. However, the cybersecurity company had more leverage to get these servers taken down than if the servers were in Russia or China, because the cybersecurity company had contacts at the major cloud computing providers in Europe. They notified these providers when they saw that their infrastructure was being used by a customer for illegal or fraudulent activity.

After learning from the cybersecurity company that the attackers were developing spoof domains for other ICSOs working in the environment, development, and humanitarian sectors, the IT team reached out to other organisations where they had personal connections with staff. They sought to warn them about the domain spoofing and ask whether they had seen any evidence of a cyberattack on their organisation. Some organisations responded by thanking them for the information, and others confirmed that they had not been attacked. It seemed like Transparency International was the only one being attacked at the time; however, it is possible that other ICSOs preferred not to share that an attack had happened or were unable to identify an attack (as few organisations have intrusion detection set up on their systems). **Transparency International has since set up an informal peer network of chief information officers to exchange information, which continues to be active**.

INTRUSION DETECTION SYSTEM

An intrusion detection system (IDS) is a system that monitors network traffic for suspicious activity and alerts when such activity is discovered.

Lutkevich (2021): ***Definition: Intrusion detection system (IDS)***

In the third week of the attack, Transparency International received another email from Microsoft to notify them that the attackers had successfully breached the phones and accounts of two staff members. The multi-factor authentication on their accounts had been bypassed through their phones. The cybersecurity company conducted a forensic analysis of the laptops of staff who experienced the breach. As staff phones were private phones, it was not possible to conduct forensic follow-up on them. At that time, the IT staff could not fathom how it could have been possible to bypass the two-factor authentication. Since learning of Pegasus spyware in 2021, they now suspect that such malware may have been used on staff phones to intercept the MFA code and transmit it to the attackers, allowing them to complete the multi-factor authentication process.

PEGASUS SPYWARE

Pegasus spyware is mobile surveillance software developed by the Israeli NSO Group and designed to infiltrate iOS and Android devices to secretly collect information. It can be installed on a target's phone without the victim needing to take any action themselves. Pegasus has extensive data collection capabilities – it can read texts and emails, monitor app usage, track location data, and access a device's microphone and camera.

Farrier (2022): ***What Is Pegasus Spyware and Is Your Phone Infected with Pegasus?***

Political figures, journalists, and human rights activists in countries around the world have been surveilled using Pegasus spyware. In July 2021, 80 journalists from 17 media organisations in 10 countries broke the story through the **Pegasus Project**. Amnesty International published this **report** and **video** sharing more insight into the spyware.

Amnesty International (2021): ***Pegasus Project: massive data leak reveals Israeli NSO group's spyware used to target activists, journalists, and political leaders globally***

The response to the attack on Transparency International was largely carried out by the IT team with support from the cybersecurity company. One staff member had some training and background in information security and recognised that they needed expert support. Before the attack, the IT team had spent a few hours per week on security, but the attack response became the majority of their work for six weeks. Although they kept senior management informed, the IT team managed the cyberattack response, as in the moment, they saw it as mainly a technical issue. In their words: "*There wasn't time to think about the big picture. Maybe we could have included the communications team to help with outreach to staff, but we didn't think about it at the time.*"

4. Outcomes & impact

The incident was contained within the organisation's systems, and there was no impact on the communities they work with or their donors. The breach did not reach any sensitive data, in part due to the organisation's systems architecture – Transparency International uses different systems to store sensitive information.

However, the attack was stressful for the IT team who were the frontline responders to the incident. They faced intense uncertainty, coming into work each day and not knowing what to expect from the attackers or how the incident might progress. However, the incident also presented a rare opportunity for the team, as it was their first time responding to a cyber incident that was actual rather than theoretical. The IT staff said, "*Now we have organisational learning from the experience, but at the time it felt like firefighting and trying to keep heads above water.*"

When a cyberattack like this occurs, staff may be less inclined to trust the operating systems they rely upon to do their work. Ultimately, the IT team and senior management at Transparency International were able to manage the situation and help staff members stay calm and vigilant and continue their work.

In the wake of the incident, Transparency International recovered quite quickly but became more conscious of security. The organisation recently launched simulations of phishing attacks to assess how many staff would fall for various methods, and they were impressed by the results. Most staff members were able to recognize most phishing attempts. This training process will be ongoing at the organisation, especially when new members join the team.

Resisting six weeks of sustained phishing attacks

5. Organisational learnings

CHALLENGES

Reporting the incident

The organisation was conflicted about whether to report the incident to the police and share it publicly. They wanted to protect other organisations by holding the perpetrators to account (although this scenario was likely unrealistic and would require years of research and forensic analysis). On the other hand, the organisation did not want to risk making themselves more of a target by going public or going to the authorities. In the end, they decided not to share the incident externally, as they wanted to return to normal operations and get back to work.

Unknowns are unsettling

Despite continuous engagement with the attackers over a period of six weeks and identification of the attackers, many unanswered questions remain. It's unclear why the attack stopped – it could be that the attackers found what they wanted, it was too difficult to find what they wanted, or they were running a time-bound campaign. Not having answers to these unknowns is unsettling for the organisation.

Balancing incident response with day-to-day demands

Dealing with the incident was incredibly time-consuming for the small IT team. At the same time, the IT team's normal work responsibilities did not go away over the six-week period of the attack, and tasks continued to pile up. Keeping up with these day-to-day demands while responding to the cyberattack was exhausting.

Resisting six weeks of sustained phishing attacks

LESSONS LEARNED

Improve threat model

Prior to this attack, Transparency International had a threat model which they used to run simulations, but it contained only vague prompts, such as "How would we respond to an attack?" The threat model has since been updated with new questions, including "Who might attack us?" and "Why would they want to do so?" The revised threat model is based on the organisation's learnings and intuition. While there are methodologies available for reflecting and creating threat models, such as the MITRE ATT&CK, these are not entirely suitable for civil society organisations. Transparency International was unable to find adequate resources, methodologies, and certifications to support the development of their revised threat model.

Establish peer networks before incidents occur

During the incident, the team reached out to IT staff at other organisations which they suspected might be targeted by domain spoofing attacks. Now an informal group on a messaging app, this peer network has become an important space to benchmark experiences and share resources. At the time of the incident, the group was of limited usefulness (as it takes time to build relationships around sensitive topics); however, if an attack were to happen again, the group would be a much stronger resource and support.

Shorten response time

With the benefit of hindsight, Transparency International thinks that their initial response was naive and too slow. It took the cybersecurity company identifying the origin of the IP addresses and telling them that the attackers were likely a state-sponsored group to make the organisation take strong action. Today, they would act faster and more decisively.

Multi-factor authentication is a useful security measure, but it is not bulletproof

Measures like using multi-factor authentication and encouraging staff to change passwords after failed login attempts helped Transparency International to keep attackers at bay for some time. However, they also learned not to consider multi-factor authentication a fail-proof measure, as it was bypassed when attackers accessed the accounts of two staff members, possibly using spyware. While there are no bulletproof solutions in cybersecurity, it's worth noting that multi-factor authentication is an extremely important tool that organisations should have in place. Without it, even the most basic phishing attacks can be successful.

Bring in cybersecurity expertise

Having an existing relationship with the cybersecurity company was extremely helpful and key to the organisation's response. By leveraging advice and support from the cybersecurity company – such as detecting the origin of the threat and conducting forensic analysis of organisational laptops – the IT team was able to fill in gaps in their understanding of the situation. The relationship existed due to a previous potential breach in the past, and Transparency International was fortunate that the cybersecurity company offered their services pro bono due to their status as a civil society organisation. However, pro bono support is rare, and civil society organisations should consider budgeting for cybersecurity expertise.

Discover more case studies
solidarityaction.network/cybersecurity

Resisting six weeks of sustained phishing attacks

Solidarity Playbook

Resisting six weeks of sustained phishing attacks

Sacha Robehmed and Nonso Jideofor

March 2023

International Civil Society Centre

CyberPeace Institute

In partnership with

Sweden Sverige

FORD FOUNDATION