WHITE PAPER

Countering the Threat of Spear-Phishing

GUIDEPOINT®
SECURITY

This whitepaper is intended as an educational tool to support an understanding of how an enterprise can take mitigating steps to minimize the risk of compromise due to spear-phishing.

Defending against phishing is essential because many resources, including the Verizon DBIR, have confirmed that between 91% and 93% of all compromises start with a phishing attack.[1]

For this paper, standard phishing or “spam” is not the focus.

Instead, spear-phishing for targeted attacks on enterprises is the focus. However, there is a significant difference between mass/spam phishing, which famously began with Nigerian princes trying to sneak money out of their country and spear-phishing that personalized the content to the intended victim.

Understanding how these spear-phishing attacks work, what the attacker process and execution look like, and finally, the vulnerabilities and mitigations of each stage are essential to defending any visible, targeted enterprise.

This paper will explore in depth the four phases of a spear-phishing attack and what steps to take to defend each step of the way.

To find out more about what spear-phishing is and how the attacks work, read our white paper: Defining the threat of Spear-phishing.

Spear-phishing Attack Defense of Each Phase

The following are the methods to defend an agency against a spear-phishing attacker, broken into four phases. The agency must mitigate the vulnerabilities in each step to thwart the attacker, as outlined below.

The four phases of a spear-phishing attack and how to defend:

1

Pre-Attack Phase

a. User training and education on how adversaries use data

b. Mandatory Compliance evaluation

c. Domain monitoring and look-alike domain verification/reporting

- Malicious actor decides to target agency
- Open source intelligence gathering
- ✓ **User training on information sharing and posting on** public facing site (including .gov)

2

Initial Attack Phase

a. DMARC in blocking mode, look-alike domain, machine learning on emails

b. Blocking of zero-day attachments and highly advanced link attacks

c. Email gateway with advanced threat intelligence and automation

- Email spoof appears to be from legitimate source
- Phishing email with malicious link or attachment
- ✓ **Technology-based security solutions that deny malicious DMARC, link and attachment attempts**

3

User Action Phase

a. User training on how to spot phishing emails

b. Phishing Reporting capabilities

- User clicks on the malicious link or attachment
- ✓ **User training on identifying suspicious and malicious phishing attempts**

4

Post-Attack Phase

a. Convert reported phish to automated active defense and remediation

b. NextGen EPP capabilities to detect malicious files and activity

c. User and Entity Behavioral Analytics to reduce time-to-discovery failure

- Endpoint protection fails to protect the system
- Credentials used to move around enterprise
- ✓ **Technology-based solutions that defend end-points and credentials from compromise**

PHASE ONE

Pre-Attack Phase Mitigation Recommendations

The best mitigation from users putting too much information in the public domain is training. Helping users understand how adversaries will use the information they put on personal social media, agency social media, or public-facing agency websites should be standard training exercises. As an example, many organizations will recommend the use of sites such as LinkedIn to keep connected with peers and members of the industry. Threat actors often will imitate a valid business person by stealing the person's photo and profile and then attempting to connect to other LinkedIn members. Once connected to an individual profile online, a threat actor can then begin to crawl through the user's business connections and company information. This tactic to steal information is part of OSINT, focusing on external data gathering.

Additionally, organizations with representatives who speak on their behalf are potential victims. As an example, the federal government allows senior ranking officials to speak at private sector conferences. In doing so, the individual reveals his or her role within an organization and information about potential projects on which the organization is working or privileged data. This is pure gold for an actor looking for spear-phishing attack vectors.

This tactic to steal information is part of open source intelligence (OSINT), focusing on external data gathering.

PHASE TWO

Defending Against an Initial Attack

There are five areas where technical solutions can identify and prevent a user from receiving a malicious spear-phishing email. Spoofed domain emails, look-alike domain emails, malicious links, malicious files, and restricting emails scored as suspicious by reputation or activity.

When attackers create imposter accounts, the most successful imposter account will utilize spoofed headers, making emails look authentic. An agency cannot stop spoofed emails without a DMARC record with full blocking mode turned on.

Codeveloped by a consortium of mailbox providers and security vendors, DMARC aims to end domain-based consumer email threats.

By leveraging existing email authentication technologies (SPF and DKIM), DMARC enables senders to instruct mailbox providers to monitor, quarantine, or reject any email that fails authentication. DMARC forensic reports will contain some or all original Headers (including the To and From email addresses), IP address of the sending email server, and either empty email bodies or full message-level data, depending on the policy of the DMARC report generator.

DMARC allows an agency to prevent someone from pretending to be an internal user fraudulently. It should be deployed in an email security stack as the last check as an email enters the agency from other entities.

DMARC.org defines DMARC as:

"DMARC, which stands for "Domain-based Message Authentication, Reporting & Conformance," is an email authentication, policy, and reporting protocol. It builds on the widely deployed SPF and DKIM protocols, adding linkage to the author ("From:") domain name, published policies for recipient handling of authentication failures, and reporting from receivers to senders, to improve and monitor protection of the domain from fraudulent email."

If an email is routed normally, the internet carriers should check DMARC as an email travels to the agency. It is possible, however, for attackers to route the email around such checks. Having a security stack that does a DMARC check itself prevents this type of highly skilled attack.

If the agency uses DMARC and is in blocking mode, the attacker will try to create an imposter pretext around look-alike domains or typo-squatting.

This is the most common type of spear-phishing seen in the wild today. If the attacker makes the look-alike domains appear internal to the organization, technical solutions can discover the domains as they are created and shut them down. Products that do this focus on brand protection, but the most significant threat would be falsifying emails for spear-phishing in the government's case.

If the attacker uses an outside look-alike domain that presents itself as a related organization with which a user would be familiar, a different control will be necessary to detect the look-alike domain. To avoid DMARC controls, look-alike threats require advanced threat scoring to mitigate illegitimate emails. It is difficult to predict what external domain the adversaries are going to attempt to use. The possibilities include every affiliated agency domain and third-party domain, such as those used by contractors or partners, which may be relevant to the individual agency.

It is unrealistic to discover and block typo-squatting of an agency or third-party domains. The best mitigation against typo-squatting is to have an email security stack that includes automated scoring for malicious and suspicious domains.

A risk scoring system can look at the domain the email is identified as the content of the email itself and any past activity associated with that address or domain, to assign a risk score to the message. If a domain has just been created, technical control can check against whois records and add a risk score. If the domain has an identifiable relationship to a reputation score from other malicious websites, it also can add to a risk score.

More advanced systems can use machine learning to score by content, dynamic reputation, and past activity (either good or bad). The system may match an email domain or address to previous activity to identify it as not risky, due to a known history related to the content. The system also can place an even higher risk score on an email due to past activity that raised suspicion, even without a full conviction. Once a score is placed on a message, thresholds for quarantine or removal can automatically protect the user.

It is still possible that a legitimate email account of a related organization, or even an internal email user previously compromised, could be used to spear-phish another user and system. If that is the case, the email may pass through with a valid payload in the form of a link or attachment, which would be the next possible threat to be mitigated. Typical malicious links and attachment security products are not designed to inherently adapt to evasive techniques that email threats allow.

Malicious links that are forwarding to a dynamic website can be sent with clean code, which alters the site information before or after a set number of hits.

This is a problem for typical link security tools that assume the link is static, check and whitelist a link that then becomes malicious later. Static site verification can be tolerable for named lists and top-level domains (TLD), but with over 100 new TLD such as .co or .biz, security vendors are pressed to validate new sites. Therefore, email link security tools that assume a dynamic link and check on click every time are required to prevent evasion.

Also, browser isolation solutions that move link clicking off-box can defend endpoints and be transparent. These solutions create temporary virtual proxies dedicated to running browsers off-box isolated from the user's system.

- Much like sandboxing, browser isolation encapsulates the user's session, cookies, and downloads into protected areas of the filesystem inaccessible by the host system. **This means that a user's sessions now are opened in a sandbox within a separate environment to completely protect the enterprise enclave and all applications on the user's system.**

When malware or file-based threats are introduced into a solution via email, a risk level must be established. Even though 'known good' or 'known bad' labels are functions of most protection suites, these suites are likened to signature-based threat analysis platforms. The deficiency with such platforms is that the only time a threat notification is triggered is when an evaluation is made upon a known threat via the detection of a hash algorithm. Advanced email attacks will leverage dynamic or polymorphic complex malware strains to utilize packing, encryption, and other methods to obfuscate their payloads, making detection and analysis more difficult.

With these threats, machine automation requires advanced diagnostics and cloud-level analysis to confirm threats. Malicious files in attachments allow for encryption that includes passkeys in the body of an email or separate email.

This type of email baiting is a threat that forces users to download the attachment and then, with a different email and password, open up the affected attachment on the host pc. Network malware sandboxes that do not have the additional context of a passkey from an email will not adapt to the evasion. A malware sandbox with integration into email security tools or an email-specific malware sandbox that can detect specific evasions is required.

As discussed prior, encrypted attachments and being sent inbound to an enterprise user should be scanned. But if they are further encrypted, and no security hash exists on the file, the file should be replaced with a message protocol policy that clearly states the message may not be safe. In cases of valid encryption requests, external users should be redirected to organization-approved, external-facing upload sites to upload secure PII information in its native format unencrypted. This will allow the enterprise to validate the security and authenticity of the user and protect the organization via compliance rules and regulations.

- It is important to note that an advanced persistent threat actor will employ spear-phishing techniques with highly evasive methods intended to avoid these technical controls, even at the cost of initial failure. The moment any attack is identified, the adversary risks alerting the SOC and organization for heightened defenses.
- This also is why reporting of suspicious emails is particularly important.

Mitigating the Risk of User Actions

Because Spear-phishing is a form of social engineering, technical controls can fail, and the attack can still be thwarted if a user does not take action desired by the attacker. However, security awareness training of users goes beyond simulated phishing to train users on what to look for to identify a phishing email and resist persuasive requests.

Recent research identifies other core messages that must be communicated clearly to users. The two most significant are that even the most protected networks can be compromised. Reporting anything suspicious in an email is more important than any embarrassment of a mistaken click or looking foolish.

Users who take the bait should be reminded that anyone can be duped, and they should report it at the first sign of a problem. It is detrimental if users do not indicate a suspicious email, even if they resisted taking the action the attacker desired. Due to the relatively high success rates of phishing and spear-phishing, a user with heightened awareness who does not fall for a phish is the best weapon against the next user clicking on the phish.

Unfortunately, most users do not report the attempt or do not know how to correctly report the attempt.

Internal metrics by GuidePoint Security assessments showed that 76% of users who correctly ignored spear-phishing attempts did not report the incident accurately. When including the 15% who did click and did not say, the total user percentage that correctly reported a phishing attempt was just over 9%.

Compliance structures such as PCI and NIST 800-53 include security awareness training that instructs users on the importance of reporting and how to report correctly. However, the success of security awareness training depends on its implementation.

EXAMPLE

Alcoa Corporation

Alcoa, an American Industrial corporation that is the world's 8th largest producer of aluminum, was targeted with Spear Phishing in 2008.

Attackers were using emails impersonating a board member and encouraged the recipients of the email to click on hyperlinks and attachments. Both of which contained malware to take over the user's computer.

The attack resulted in the loss of thousands of emails and attachments containing proprietary information.

The best template for a security awareness training platform will allow for:

- **Streamlined communication to the end with real-time reporting on the awareness process**
- **Knowledge assessments**
- **Simulated attacks**
- **Interactive training modules**
- **Customizable content**
- **Automated analysis of reported phish suspects**
- **Pre-built integration with the email security stack**

There are other features to consider, but each of these should be high on any agency's list of requirements to ensure a successful program.

PHASE FOUR

Ensuring a Post-Attack Defense

After the user is duped into taking the specified action of a spear-phishing attack, it is no longer a prevention question at an email level. It is about defending the network and finding and removing the adversary. A targeted phishing campaign could be reported with adequate training by a user within minutes of another user taking the bait. A valid email security stack and an automated assessment that initiates quick action are essential.

Most non-spear-phishing attack reports should be filtered to reduce the noise that the SOC will have to deal with. Without a valid email security stack, a SOC will have a solid business case to procure the required technology and eliminate the deluge of spam phishing email reporting with practical security awareness training. The advantage of solid user education and compliance validation on phishing prevents the 'duped user' effect of 'click first and ask questions later.'

Therefore, building an effective malware reporting practice into its onboarding program is essential to the organization's security.

Once end-users are trained, and reporting is reduced to (primarily) advanced phishing and spear-phishing, the next important factor is time-to-identification. The automated assessment function should significantly reduce the workload and time required by the SOC to develop a reported phish into an actionable SOC playbook.

After confirming an attack, the most important next action is tracking down every user who received the phish and following up with remediation and detection if that user could possibly have taken the bait. Integration with the email security stack and a SOAR (Security Orchestration and Automated Remediation) can significantly help the SOC.

Orchestration tools assist administrators through a process of machine automation in identifying the threat, reporting the threat within a ticketing system, and providing an automated remediation action to endpoint systems, such as temporarily quarantining off the network any system that received the threat without reporting it. Other examples of this include quarantining all emails forwarded post-delivery to users and automatically enrolling users into additional compliance and phishing training.

Users who repeatedly fail to report attempted phishing will be motivated to inform the next one if an EDR quarantines their system for analysis because they did not report it.

The next post-attack step assumes that no one reported the phish, probably because it was targeted so precisely that only one user who took the bait received it. The next defense is at the endpoint with a complete EPP that includes Next-Gen AV and EDR capabilities.

If the EPP is unable to prevent or detect the attacker from taking over the user's system and possibly gaining the user's credentials, the next step in defense is identifying a system (called an "entity") or a user credential that has been compromised. A compromised entity or credential always will act differently than the regular user and entity when used for malicious purposes. A UEBA (User and Entity Behavior Analytics) can quickly identify the entity or user by changing behavior and reducing the time-to-identification.

Gartner defines an EPP[2] as follows:

"An EPP is a solution deployed on endpoint devices to prevent file-based malware, to detect and block malicious activity from trusted and untrusted applications, and to provide the investigation and remediation capabilities needed to dynamically respond to security incidents and alerts."

Conclusion

Spear-phishing is the most difficult type of attack to defend. An adversary looking to exploit users can deploy it quickly to thwart an organization or user's defenses. As this paper has detailed, an email security stack explicitly designed to detect and defend against a spear-phishing attack is essential. Still, it is not the only component of an effective solution. Building a people-centric view of the enterprise that supports a comprehensive cybersecurity program is any organization's goal.

Data Analytics, link detection, advanced attachment analysis, DMARC, domain spoofing in several forms, web isolation, and end-user security awareness training are all part of a complete solution. Finally, interconnecting individual point products by integrating all network defense phases that leverages machine learning and advanced threat automation for mitigation is the most effective path for any agency's email security.

Endnotes

1 – Dark Reading "91% of cyber attacks start with phishing":
https://www.darkreading.com/endpoint/91--of-cyberattacks-start-with-a-phishing-email/d/d-id/1327704

2 – Gartner EPP MQ:
https://www.gartner.com/doc/3848470?ref=mrktg-srch

GUIDEPOINT®

SECURITY

2201 Cooperative Way, Suite 225, Herndon, VA 20171
guidepointsecurity.com • info@guidepointsecurity.com • (877) 889-0132

WP-cSPEARPHISHING-082020-02