AVERTIUM

LogRhythm

4 Trending Phishing Techniques and Tips for Detection

Table of Contents

4 Trending Phishing Techniques and Tips for Detection

Introduction

No other cyberthreat plagues the world like phishing attacks. This email-borne threat action tops the list as the most used initial attack vector in data breaches today[1]. What began decades ago as simple spam intent on tricking recipients into visiting sites and becoming customers has grown into a world-wide industry. Based on solid social engineering principles, these email attacks seek to deceive recipients into engaging with malicious content, commit fraud, or perform desired business actions based on the scam at hand.

According to the FBI, phishing attacks are the number one internet crime by victim count[2] and an overwhelming majority of organizations (88 percent) cite having experienced phishing attacks[3]. This is a problem that shows no signs of slowing down.

In the end, phishing attacks can have a significant impact on organizations, with the most common including loss of data, credential compromise, ransomware infection, other types of malware infections, and financial loss. Lots of security solutions aim to stop phishing attacks before they reach the inbox, but because of the efforts by cybercriminals to evolve their tactics, it should be assumed that some percentage will always get through.

What's needed then is a concentrated effort in trying to strengthen the weakest point in your security strategy: your users—the human factor in the equation. And, should your users fail the organization (spoiler alert: they do), it's important for security teams to have an ability to detect phishing attacks themselves.

In this white paper, we'll dive into the makeup of phishing attacks, take a look at what aspects of emails users need to focus on to elevate their security vigilance, and discuss how you can monitor some of those same characteristics to detect phishing emails—even when security solutions designed to do so don't.

[1] Verizon, Data Breach Investigations Report (2019)

[2] FBI, Internet Crime Report (2020)

[3] Proofpoint, State of the Phish Report (2020)

Detecting and Investigating Cyberthreats with LogRhythm

No security solution provides 100 percent protection against phishing attacks. So it's critical to have centralized visibility of all activity and changes within your entire environment to understand when and how attacks are occurring.

Look for insights in this paper from LogRhythm and examples of how their NextGen SIEM and NetMon solutions help to detect and provide detail around phishing attacks.

Phishing 101: Phacts About Phishing

Before we can dive into how to identify a phishing email and how to best prepare users to play a role in your security strategy, we first need to level set some of the what, how, and why of phishing.

At its core, phishing is an email-borne attack designed to get the recipient of the email to act in a desired way—whether it's clicking a link, opening an attachment, giving up information in a reply, or performing a business-related action (e.g., initiating a wire transfer). It's important to note that, in almost all cases, the attack is solely dependent upon the recipient to engage with the email's content. So, phishing attacks use social engineering—the art of deception used to manipulate the victim. The techniques used accomplish a few things:

- **Establish credibility:** It's important to lower the recipient's defenses. The moment they smell a rat, the scam is finished. One method is the use of impersonation, which occurs in over two thirds (68 percent) of phishing attacks[4]—where the scammer uses a lookalike domain, fakes the display name of a brand or individual, or uses a compromised account. Another method is using SSL as part of the connection to a fake website. It makes the recipient believe the site is legitimate (that is, unless they look at the certificate and find it's a spoofed domain name). One last method used in attacks intent on compromising credentials is the spoofing of website logon pages to provide victims with a familiar interface to not raise suspicion.
- **Create urgency:** A cybercriminal needs the victim to engage with their email the moment it's opened; otherwise the likelihood of both fooling the victim and getting them to take the desired action reduces significantly. Urgency is generated by building an emotional connection with the email and its message. Anything from "you've won!" messages to "there's a problem," to the victim's "boss" telling them to do something ASAP is used by phishing scammers.
- **Avoid detection:** The idea of simply attaching a malicious executable as the attachment in an email is dead. This legacy method of infection has long been detectable by security solutions. Instead, cybercriminals use a number of steps, using urgency and impersonation along the way, to get users to follow what may be two or three steps to finally release the malicious code onto their endpoint or enter in their online credentials. These additional steps are taken specifically to keep from being detected by security solutions.

[4] Agari, Email Fraud & Identity Deception Trends (Q1 2020)

LogRhythm Insights: Impersonation and the CEO

No one has more credibility and can create more urgency than the CEO of an organization. It's why scammers do their homework and target organizations with emails impersonating the CEO.

Using nothing more than social engineering, these campaigns take the simplest approach: the CEO needs your help. Take the example of the extremely basic email to the right:

In many cases, this is all it takes to establish that it's the CEO asking for help. Because these tend to use external email addresses, the only telltale sign is a complete mismatch of the email address used and the person they are impersonating.

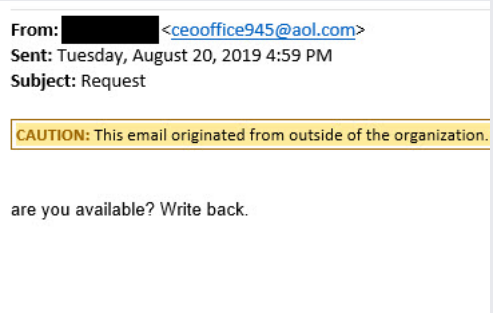

Figure 1. Example of social engineering phishing attempt

Phishing Goals

A phishing email is just the first step in a chain of malicious events. Let's look at the various types of initial malicious content found in phishing emails and cover the possible end game goals for each of these phishing attacks.

- **Malicious attachments:** Typically, these attachments contain content used to obfuscate their intent and avoid detection by security solutions, all while working to communicate with a command and control server and, eventually, download the malicious payload to be installed. The common goals of malicious attachments are to install malware used to provide remote access to the victim network, install ransomware, steal data, further proliferate malware infection by sending emails on behalf of the logged-on user, or leverage access gained to gather intel, island hop, or commit fraud.
- **Malicious links:** Links can provide the same result as attachments (through the installation of malware), as well as use scams designed to compromise online credentials to Office 365, accounting platforms, and other cloud-based applications.
- **Malware-less emails:** Some phishing campaigns rely purely on social engineering and use no actively malicious content. Business email compromises (BECs) and the previously mentioned CEO fraud often take this approach with the intent of convincing the victim to take actions such as modifying banking details, wiring money, purchasing and sharing gift cards, and providing internal company details.

LogRhythm Insights: Credential Theft—an Attack Without Malware

One of the easiest ways for a phishing attack to be spotted is by an attachment. It's one of the reasons some cybercriminals have chosen to take the path of using phishing to steal online credentials. The attacker doesn't need malware—just some crafty social engineering and believable spoofed logon pages.

For example, take a look at the logon page below. The scam involved notifying a user via email that they had an unread document in the cloud they needed to look at. The email provided a link, which brought them to the following page:

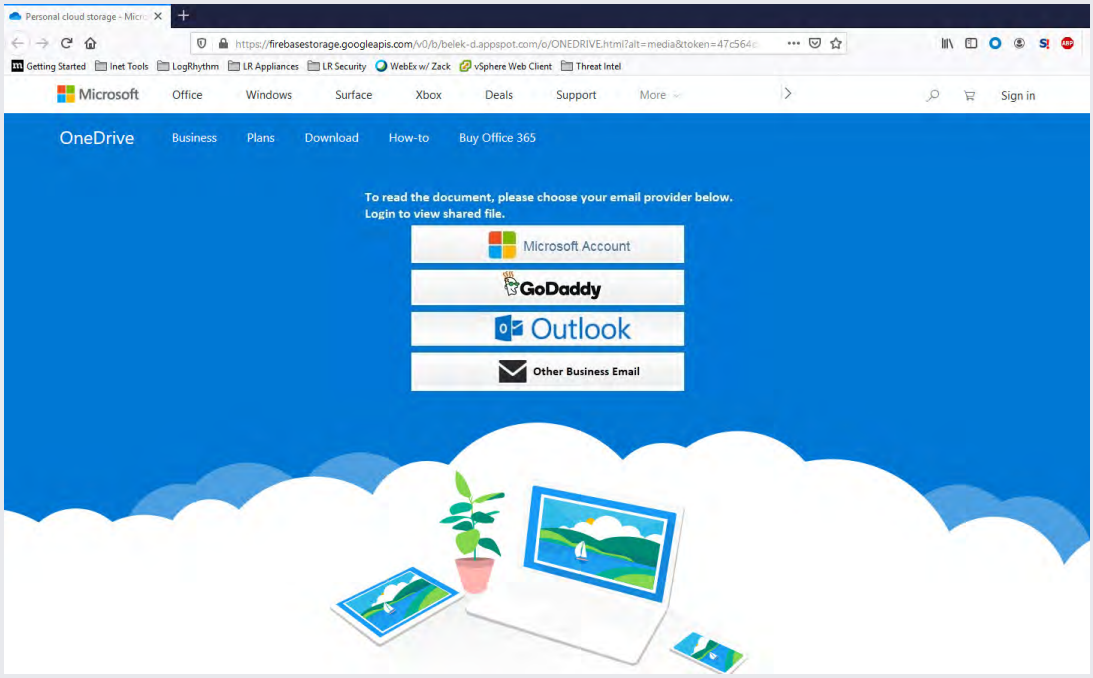

Figure 2. Example of a scam logon page

While it looks legitimate, if you note the URL, it's being hosted on Google's Firebase storage. The logon form used in this scam (which is presented upon clicking one of the buttons) is actually hosted on formsite.com. The unwitting victim provides their credentials to see the document and is instead presented with an error message that the file couldn't be opened.

If executed perfectly, the user may not even suspect they're a victim. So, it's necessary that organizations monitor for this kind of activity. LogRhythm's NetMon can help provide needed insight into suspicious network activity.

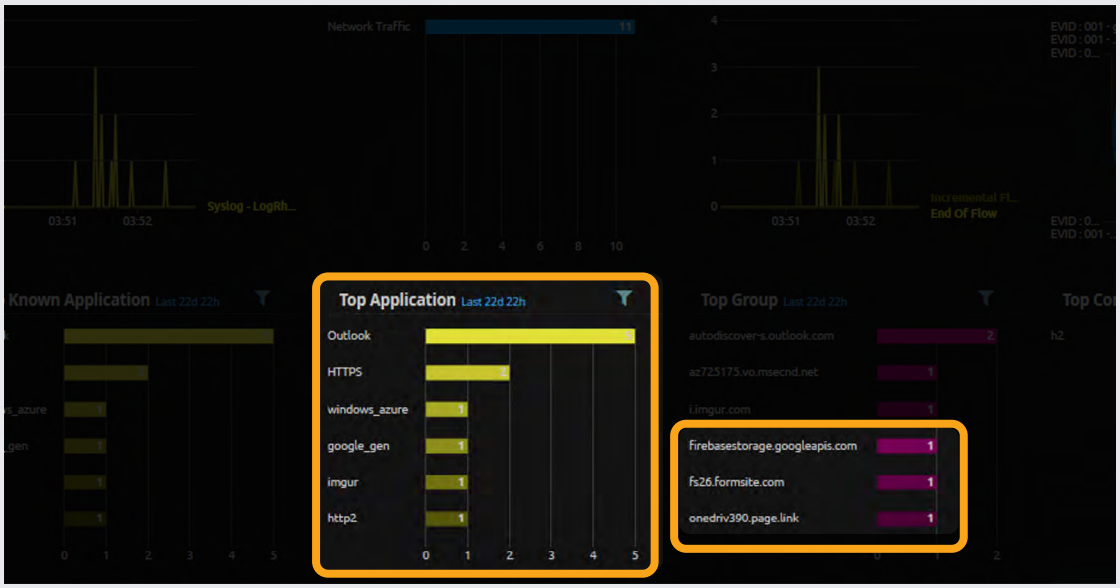

Figure 3. LogRhythm NetMon correlates traffic sources and web traffic to detect abnormal patterns that may indicate an attack

With tools like NetMon, organizations can quickly detect abnormal traffic patterns that may indicate an attack—such as by correlating Outlook application traffic with formsite.com web traffic—as well as ascertain the scope of any successful attacks for remediation.

As you can see, regardless of whether malicious attachments or links are used, social engineering plays a significant role in a phishing attack in order to fool the one part of your organization's security that isn't a piece of software—your users.

It's important to look at your users as a human attack surface and look for ways to strengthen this insecure part of your security stance.

Hardening the Human Attack Surface

This is not a simple task, as it's not reasonable to think that users will respond to varying phishing attack methods consistently, like a security solution would. But there is the opportunity to educate users on how phishing content differs from legitimate emails. The trick here is to get users to begin to think about email content they engage with in more of a technical aspect, rather than a social one.

For example, if an email comes to someone in the organization's finance department stating there's been a problem with a payment transaction, the social side of the user reads the email, involves their emotion (in this case surprise, concern, etc.), and drives them to click on the malicious content. But if the user reviews the email from a technical perspective, that can deter them from clicking on the malicious content and focus on determining whether a given email is valid or not.

> *So, what parts of an email help identify that it's malicious?*

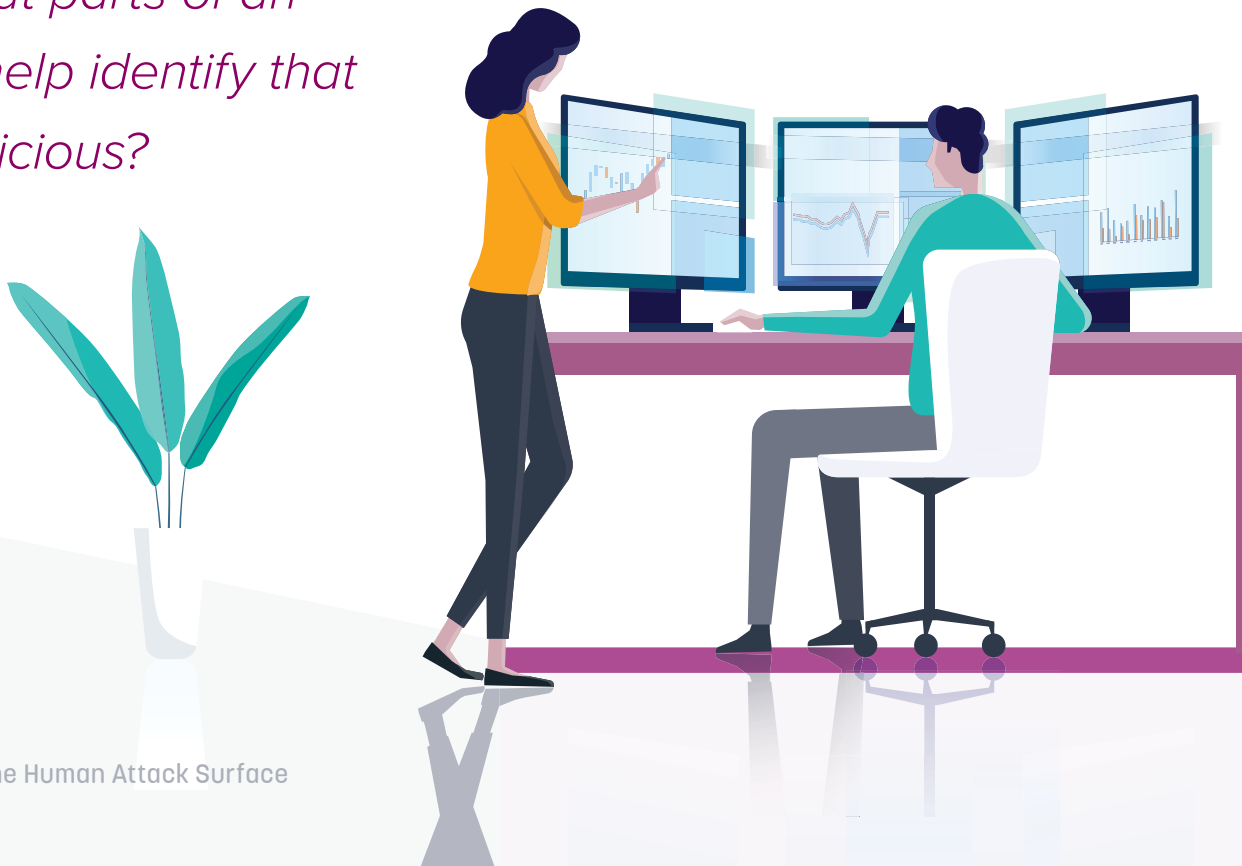

Email Elements to be Scrutinized

As IT professionals, we can fairly easily tell if an email is legitimate or not. But cybercriminals are getting better at their craft, making emails and web pages look, sound, and feel increasingly legitimate. So, it's important to be thinking about all the various ways users and IT can detect a phishing email.

Following is a non-exhaustive list of the various elements that attackers might use in a phishing campaign, and what both users and IT should be looking for to identify an email as suspicious, if not malicious:

Sender/sending details: Start with who is sending the email in the first place. Look at the domain the email is purporting to be sent from, looking at the spelling, use of homographic characters (lėttèrs that look like those in the actual domain) to impersonate a company or an individual. Also, take note of the email address and name of the sender. Is this sender someone that frequently exchanges emails with the recipient? Is it claiming to be your CEO, but is coming from a gmail.com domain? The alignment of sender details is a good first indicator that something may be wrong. IT and security teams can additionally look at the IP address of the server sending the email, the age of the domain, DNS servers, domain registrar, and SSL certificate authorities as ways of validating authenticity.

Recipient: Is the recipient of the email in a higher-risk category, such as someone with access to financial information, intellectual property, customer data, etc.?

Subject: Usually more akin to spam detection, looking at the subject can still help determine legitimacy. Look for misspellings, incorrect grammar, and any other signs that the email is unusual or abnormal from those emails usually received.

Body content type: While most emails are HTML these days, it's important to note whether the email supports tags and links that are used commonly in phishing emails.

Links: Initially, users can usually hover over a link and see where it points to. That's a good first step. IT and security teams may want to follow the link to its final resolved URL, IP address, and content delivered to determine whether it's malicious.

Attachments: Just the presence of an attachment from someone the recipient doesn't know is suspicious. Additionally, the type of attachment makes a difference as well. For example, receiving a password-protected Word document from someone you regularly do business with (that has never sent one of these before) is suspicious. Both need to be considered. Lastly, consider the content itself. If that Word document is a "proposal" that you're not expecting (nor do you ever receive proposals from the sender), it should be deemed suspicious.

LogRhythm Insights: Legitimacy Should be Found in Every Detail

Many phishing scams leverage existing trusted file hosting sites—such as Dropbox, OneDrive, and Google Drive—as locations to store their malicious files. However, this is obfuscated to the recipient through social engineering methods where the file's location is used to legitimize the email. Take the example below that comes from Dropbox:

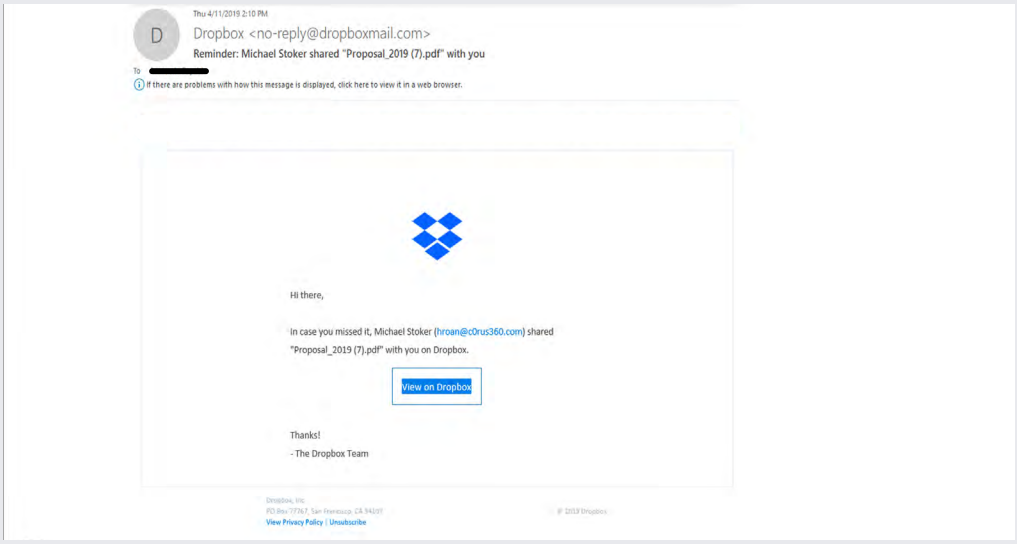

Figure 4. Example of an attack utilizing a file-sharing site to store malicious files

At first glance, this looks simply like a courtesy notification from Dropbox. The sending domain is legitimate, the link actually points to dropbox.com, and the person "sharing" this PDF is presumably someone known to the recipient. Upon further inspection, there is one indication that this may be a phish: **the email domain indicated by the sharing party is @c0rus360.com**. But the PDF link points to Dropbox! Surely this is legitimate. As it turns out, the PDF is stored on Dropbox, but the linked PDF file contains a URL to an unsafe location—in this case, to a spoofed Office 365 logon page attempting to steal credentials.

It's important for users and IT pros alike to understand that if any one of the elements looks suspicious, it is a good indicator that the whole thing is bogus. Remember that odd domain c0rus360.com? Turns out it was registered last year, while the actual corus360.com was registered in 2008. Cybercriminals will create, use, and dispose of domains on an as-needed basis, making the domain creation date a key factor for IT and security teams in determining email legitimacy.

But Can You Count on Users?

It's important to include users in the organization's security stance. But convincing them to think about security, when their job is marketing, for example, is no easy task. There are a few ways to accomplish this.

One is through security awareness training, where you teach users the importance of them practicing secure habits, as well as how to spot the elements mentioned above. Users should also be taught to balance determining suspicion with legitimacy. Without doing so, your organization will end up with a lot of false positives.

You can also take a more active approach by periodically phishing your users. Phishing testing provides IT and security teams with a feedback loop on where their security (read; which user or users) is weakest. Testing also helps to reinforce the security culture of the organization.

Even so, it's critical for IT and security teams to understand the risk of placing such a burden on the user and the repercussions should the user fall short in their newfound "duties." On one hand, 39 percent of employees surveyed don't even know what phishing is![3] So, it's obvious that putting some form of security awareness training in place makes sense. But on the other hand, it's estimated that 60 percent of users reported phishing emails were actually false positives[4]. So, the overzealous security-minded user may not be as adept at spotting a truly malicious email as the seasoned IT pro.

What's still needed is a way for IT and security teams to provide their own means of detection, should protective security measures—both technology and human—fail.

Adding in Another Layer of Detection

Because security solutions are designed to detect specific aspects of phishing, and because humans as a detective-form of security are going to be fallible, it's important to have visibility across all parts of your environment—endpoints, networks, security solutions, firewalls, and more.

LogRhythm Insights: Using Multiple Sources to Detect Phishing

Detecting a phishing email takes more than just scrutiny; it often requires a deeper look at the series of actions being taken before it's recognized as being malicious. The activity created by the simple clicking of a malicious attachment or link may only be partially recognized by a given security solution. What may be needed is an ability to centralize and review disparate data from a variety of network environment sources and security solutions to understand whether suspicious activity is actually malicious.

LogRhythm supports any data source that can export to Syslog and can consolidate multiple data sources into its powerful SIEM. In the case shown below, an endpoint protection solution can be used to log when Outlook spawns a browser as a child process, passing this data to LogRhythm. Once consolidated within LogRhythm, this and all related activity data can be monitored for and alerted on.

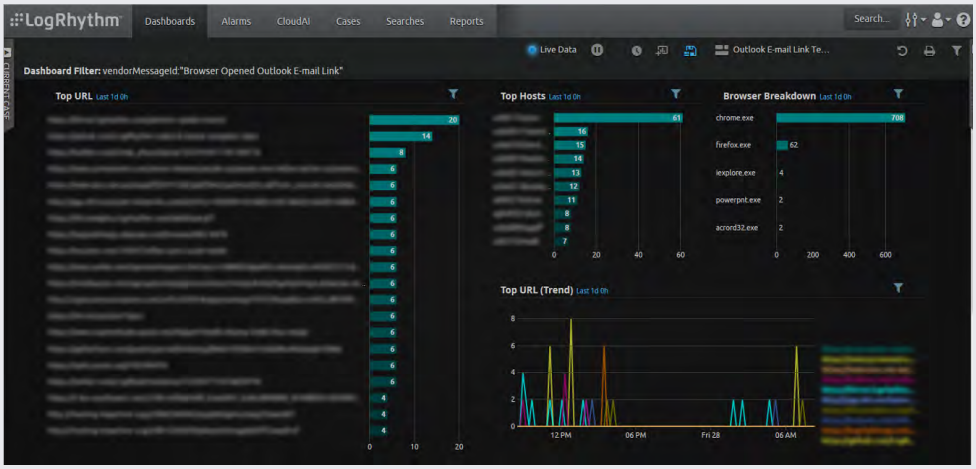

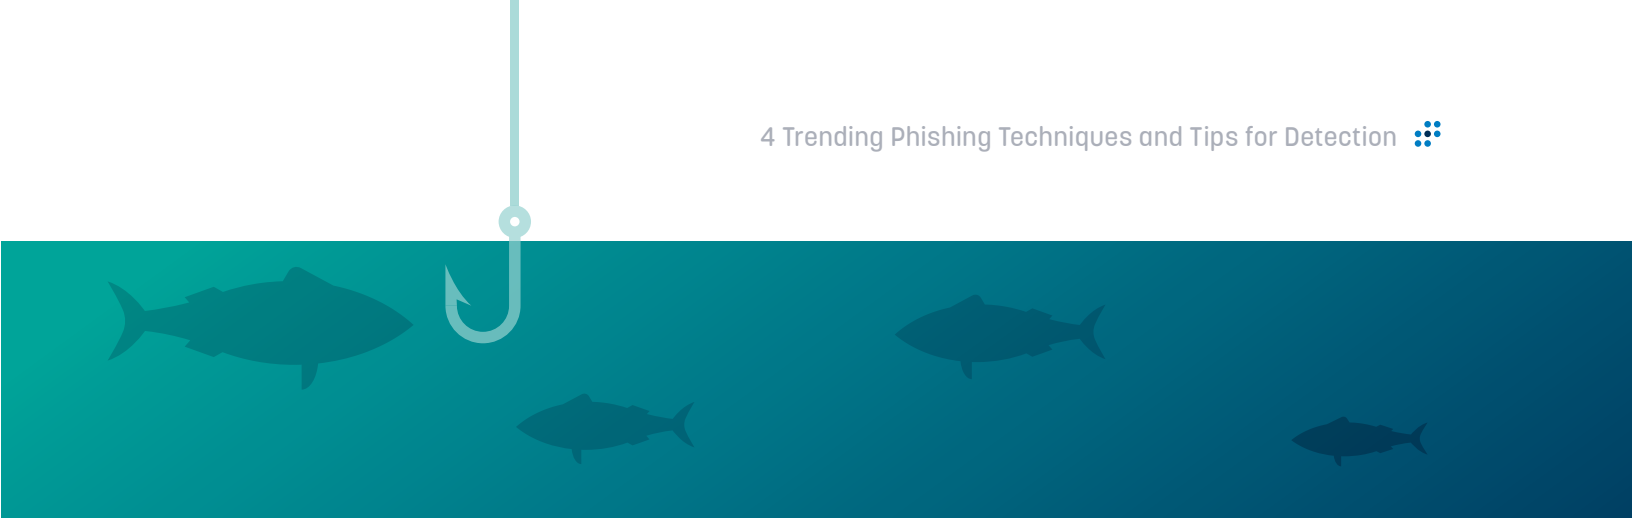

In essence, from a system's point of view, if you were to centralize pertinent security and activity data, how do you configure rules to recognize signs of a phishing attack?

There are several data sources to choose from—some of which have been covered earlier in this paper. Following are some more signs to consider that can be a part of your additional layers of detection. Most require integration with specific data sources and, therefore, a more comprehensive strategic view on how you'll monitor for malicious email.

- **Email details:** Something as simple as the sending domain and the sending server's IP address can be compared against the sender policy framework (SPF) and DomainKeys Identified Mail (DKIM) records, as well as DMARC policy records within DNS. Sending domains can also be compared with reputation data sources.
- **Attachments:** A given attachment can be compared with known malicious file types, as well as with any analysis output from antivirus, email scanning, and endpoint protection solutions.
- **Links:** Because the link provided in a phishing email is often not the final destination, network traffic data can be used to perform reputation checks for each step along the way, along with reviews of all SSL certificates.
- **Sender:** Impersonation attacks attempt to make it appear that the sender is known and/or trusted. With an ability to scan emails, it is prudent to run checks of the display name against the name portion of the email address, internal users (if the email was supposedly internally sent), and the reply-to address (they should match). Additionally, review domain names for look-alike domains using character switching, homoglyphs, and homographs, as well as long domain strings to obfuscate the real domain being used.

Detecting Phishing Attacks: One Way or Another

Phishing isn't going anywhere anytime soon, as long as cybercriminals have success with getting unsuspecting users to open malicious attachments, click on malicious links, and perform the needed action to give the attacker exactly what they want. So, it's imperative that organizations have an ability to detect and investigate phishing attacks within their environment.

Security solutions can provide a viable layered defense against malicious emails that includes scanning email content before it reaches the inbox. But as attackers evolve their attack techniques, it should be expected that some malicious emails will get through. Because of this, the next layer of defense that needs to be propped up is the user. It should be assumed that there will be varying levels of success—even with security awareness training in place.

So, the last layer of defense is visibility. Without visibility, it's sort of like making the joke statement: If a malicious email steals your credentials, but you don't know it, does it make a sound? There are countless potentially malicious actions that occur on your network daily. The trick is to have central visibility into every network request, every email, every process spawned, and more to be certain that no phishing attack goes unnoticed.

The goal is to not let a single phishing attack get by—or, at very least, not go unnoticed. So, having a layered strategy in place that is built on detection, hardening the human factor, and complete visibility will minimize the risk of successful phishing attacks while improving your organization's ability to detect and remediate them.

About the Authors

Randy Franklin Smith is an internationally recognized expert on the security and control of Windows and AD security. Randy publishes www.UltimateWindowsSecurity.com and wrote The Windows Server 2008 Security Log Revealed—the only book devoted to the Windows security log. Randy is the creator of LOGbinder software, which makes cryptic application logs understandable and available to log-management and SIEM solutions. As a Certified Information Systems Auditor, Randy performs security reviews for clients ranging from small, privately held firms to Fortune 500 companies, national, and international organizations. Randy is also a Microsoft Security Most Valuable Professional.

Brian Coulson, Threat Research Senior Engineer in LogRhythm Labs, works to keep abreast of current cyberthreats and news, develop threat detection and response content, and demonstrate how we detect and respond to threats. In this role, he regularly engages with the LogRhythm Community and offers advice and solutions to remediate common security-related issues. He is also responsible for creating new content in the form of AI Engine rules, WebUI dashboards, and Kibana visuals.

Eric Brown is a Senior Security Analyst and part of the Office of the Chief Information Security Officer team, where he helps maintain LogRhythm's overall security posture. He handles incident response, conducts threat hunting, detects and analyzes anomalies, analyzes and replies to reported suspect and phishing emails, and monitors logs and other network traffic analysis within the LogRhythm SIEM. He is also part of the Change Advisory Board (CAB).

About Avertium

Avertium is one of the largest cybersecurity services providers to the mid-to-enterprise market. Forged out of three award-winning cybersecurity services companies, each with a unique perspective on the security landscape, Avertium brings enterprise-level security to the many mid-sized and larger organizations that don't have access to comprehensive, specialized protection. Avertium is a LogRhythm **Service Authorized Partner** with experience in deployment, management and configuration, etc.

More than 1,200 organizations in industries ranging from financial services and manufacturing, to technology and healthcare benefit from Avertium's managed security, consulting and compliance services delivered with more rigor, more relevance, and more responsiveness. The company's dual security operations centers are located in Arizona and Tennessee.

Avertium. Show No Weakness.™

avertium.com // hello@avertium.com // 1.877.707.7997

1.866.384.0713 // info@logrhythm.com // 4780 Pearl East Circle, Boulder CO, 80301