proofpoint

Longline Phishing:

Email-borne Threats, Cloud Computing, Big Data, and the Rise of Industrial Phishing Attacks

A Proofpoint White Paper

Contents

Introduction

The last few years have seen a dramatic increase in the use of email as a vehicle for cyberattacks on organizations and large corporations. Such attacks have evolved from the simple inclusion of malware as a non-disguised executable file, to more socially engineered "phishing" style attacks, which attempt to persuade the recipient to activate malware or voluntarily provide valid security credentials—often simply by clicking a link that leads to a malicious or fraudulent website.

Phishing in turn has evolved – from spam-like mass emails ("You've won the lottery, click here") to "spear phishing" (a phishing attack that uses customization methods superior to those used in mass phishing attacks, such as targeting delivery service users with "You have a tracking problem, click here"). Spear phishing has further evolved to today's "Advanced Targeted Attacks" which leverage a variety of sophisticated techniques—sending email campaigns with extremely low-volume, specifically targeting a single organization (or even a single individual or small group within the target organization), and using a highly-customized message, carefully crafted after researching the target.

Until recently, email defense systems have been fortunate in that attackers faced a cost/volume trade-off. That is, crafting an email-borne attack that was highly unique and highly randomized (and thus more likely to pass defense systems) was a largely manual effort, which limited the scope of such customized attacks. Attacks that were more broadly-distributed were less customized and more easily filtered by email security solutions, as a result. Both types of attacks resulted in sufficiently low penetration rates that IT teams often had a chance to detect and remediate such breaches before significant harm occurred.

However, today's advanced phishing tactics may have overcome the cost/volume trade-off. Borrowing tactics from cloud computing and database marketing, attackers are now engaging in industrial-scale phishing attacks that leverage sophisticated customization and delivery techniques.

Proofpoint has dubbed these "Longline phishing attacks" or "Longlining" after a common industrial-scale, commercial fishing technique. Proofpoint researchers have found that these phishing attacks have markedly higher penetration rates than traditional attacks. They also have surprisingly high recipient clickthrough rates—higher than 10% in the attacks that Proofpoint researchers studied for this report.

Longlining Defined

Longline phishing attacks are distinguished by three specific characteristics:

1. **Proportionally low volume per organization, with high overall volume.** While not as targeted as an Advanced Targeted Attack (Longlining often hits tens of companies simultaneously), the volume of email per attack received by individual organizations represented far less than 0.1% of their overall mail flow. Across all targeted companies, however, a Longlining attack will likely send tens to hundreds of thousands of email messages in a few hours.

2. **Aggressive obfuscation and customization techniques, including:**
   - Massively rotated sending IPs and spoofed sending addresses
   - Malware hosted on dozens of compromised sites
   - **Text customization,** ranging from minor wording shifts for "hashbusting" purposes to significant title and body content changes based on sending time and recipient company
   - **URL rotation and/or link obfuscation**, where links are typically obfuscated in HTML and may also be shortened and/or made unique via shortening techniques
3. **Malware payloads that leverage zero–day exploits.** Links contained in Longline phishing messages lead to malware that exploits security holes for which no patch has yet been released, or for recently–discovered security holes that are likely to be as yet unpatched in most organizations.

The net result is that individual messages received are largely unique, and thus successful in bypassing traditional email security systems. No organization hit by a Longlining attack will receive more than a few email messages with the same characteristics.

In addition, the senders, recipients, and embedded URLs will appear valid and reputationally positive, making the emails very difficult to detect via conventional methods—even though the total volume of messages sent in a given attack may spike into the hundreds–of–thousands when considered globally.

Preparation of Longline Attacks

Longlining in the physical world of aquatic fishing requires equipment, in the form of a platform (a boat), bait, and hooks and lines. Similarly, in the electronic world of phishing, Longlining also requires equipment in the form of a platform (botnet or snowshoe network), bait (components of an attractive email), and hooks and lines (compromised sites and malware).

The Platform: Botnet/Snowshoe Networks

Even as legitimate marketers are harnessing the power of big data and cloud computing, attackers have adopted similar techniques. While it's rare to see an attacker rent computing space on a public cloud, botnets (a network of compromised computers under an attacker's control) and snowshoe networks (a botnet deliberately spread across a wide range of disparate IP addresses) are readily available to rent in black market forums.

The result is that attackers have computing power at their disposal which can far exceed the datacenter capacity of any single target company. These can be used for simple dispatch of email, or for more sophisticated data processing, such as multidimensional "mail merge" activities, and combining a rotating set of URLs, senders, text and recipients to create phishing attacks in which every email message is essentially unique. In one Longline attack observed by Proofpoint, nicknamed Letter.htm after the destination URL ending, more than 25,000 different sender IPs were used.

The Bait: Email Components

Just as legitimate marketers can purchase organizational charts and contact lists of names, titles, and email addresses, attackers can also buy (or obtain) such lists on the black market or even via legitimate means.

For example, knowing that a target list contains IT purchasers of a certain software or technology, can help attackers craft a significantly more compelling set of rotating text "lures" and URLs used in the email. As with legitimate email marketing efforts, the attackers' goal is to compel the recipient to click a link. In the Letter.htm attack, Proofpoint observed more than 35,000 sender email address aliases rotated with different IP and URLs through 185,000 emails.

Hooks and Lines: Compromised Sites and Malware

To disguise the attacks from both email filters and end-users, the URLs used in the emails will ideally lead to known websites with a positive reputation. It's far more likely that an end-user will click on a link to a site that they think they know, rather than an anonymous, out-of-country site.

Correspondingly, attackers will often compromise legitimate websites, gaining access to the sites but waiting to load malware (or a browser redirect to malware)—on a path deep within the site—until after the attack has launched. In the Letter.htm attack, Proofpoint observed twenty-two compromised websites, with the malicious payload at an average path depth of three subdirectories.

Similarly, the malware used in such attacks has also evolved. Many browser exploits are effectively invisible, downloading and executing code in the background while a web page appears to still be loading (these are often termed "drive-by downloads"). Often a user will be completely unaware that they have become the victim of a successful attack, as they'll experience nothing more than a slowly loading site, or at worst a web page timeout.

Common exploit *detection* kits can unfortunately also be used as *exploit* kits, providing attackers with a point and click "build a malware package" interface. In the Letter.htm attack, a JavaScript "Blackhole.js" variant was used, capable of affecting most major browsers.

Execution of Longline Phishing Attacks

Armed with a botnet, set of target lists, access to destination websites with positive reputations, and a malware package, an attacker is ready to start their Longlining trawling run.

Initial Probe and Testing

Attacks that Proofpoint has observed are typically delivered in a classic bell-curve progression. An initial small set of rotated emails will be sent, effectively testing the defenses of the various targets. If certain URLs or company defenses are already detecting the attack, it makes sense to take that feedback and eliminate those choices from the variables available to the attacker's email generation system—the mail-merge "baiting" the hooks.

Figure 1 on the next page shows the email distribution pattern for one of the Longline phishing attacks ("Mail.htm") observed by Proofpoint during the course of this research.

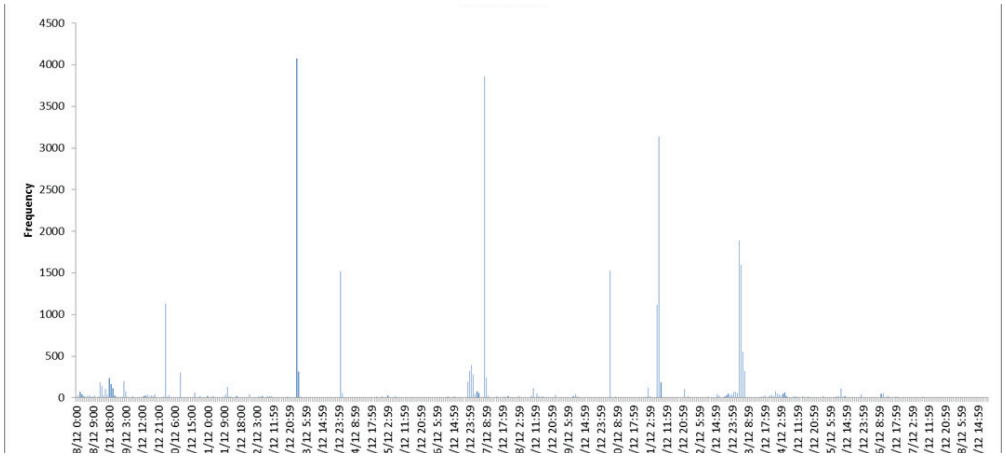

Figure 1: Mail delivery pattern for one of the Longline phishing attacks observed by Proofpoint. The early stages of the attack (left hand side of the chart) show the initial testing period, where a small number of messages are sent. After the initial period of testing, the attack becomes more aggressive. In this case (dubbed "Mail.htm" by Proofpoint researchers), phishing emails were sent in one-hour bursts over a two week period.. Proofpoint observed 28,800 email messages in this attack, sent from 2181 unique IP addresses, delivering lures that led to 803 unique malware-infected URLs.

Volume Delivery

Having established a successful set of variables, attackers then deliver a high total volume—but low proportional volume—of email in a relatively short period. The Letter.htm attack observed by Proofpoint took about three hours, including initial probing, and ultimately delivered more than 185,000 emails to 80 companies. In each case, the volume of email received by each organization represented less than 0.06% of its total daily mailflow, ensuring that the attack was effectively stealthy on a per-receiving-organization basis. No single company received more than three emails with the same variable set. In many cases, the content of each message received by a single organization was entirely unique.

Long Tail

Often a successful attack will generate additional successful attacks, as the infiltrated malware yields additional names and targets. In the observed Letter.htm attack, a second wave attack was initiated approximately three weeks after the first attack—possibly timed to reinfect organizations that had remediated the effects of the original attack.

Results and Effectiveness of Longline Phishing Attacks

While the mechanisms used in the delivery of Longline phishing attacks are worth understanding, they would be of only academic interest if they were not also *highly effective* in luring message recipients to take the desired action. Several key findings from Proofpoint researchers are reported below.

For the following analysis, Proofpoint observed more than a billion email messages, delivered over several week-long periods to multiple "Fortune 1000" enterprises.

Proofpoint's research team observed that:

- **More than a quarter (27%) of email messages classified as spam also contained links to malicious URLs.**

- Of the Longline phishing attack messages that escaped detection by traditional perimeter defenses (i.e., messages that were successfully delivered to a legitimate recipient inbox), recipients fell for the attacks at an alarming rate.

  More than one in ten (11%) recipients clicked on embedded links to malicious URLs, effectively inviting attackers into their organizations.

  The majority of such URLs link to exploits that are undetectable to an observer (e.g., the exploit looks like a web page or browser waiting to load).
- Mobile and remote users were dramatically more susceptible to Longline phishing attacks. **Nearly one of every five (19%) clicks on malicious URLs embedded in email occurred "off network"**—that is, outside of corporate perimeter protection—when employees accessed their email from home, on the road, or via mobile devices.
- **Approximately one in seven (14%)** malicious URLs are sent only to a single, targeted organization. On average, most URLs aren't sent to more than five organizations, making the malicious URLs very difficult to detect and stop using conventional, signature-based methods.

In short, Longline phishing attacks are not only effective, but are also designed to circumvent existing perimeter-only security systems. Given the frequency of off-network clicks and unique URLs/malware signatures, Longlining would seem to necessitate a "follow the email" protection system to ensure that users are not compromised when they are outside of the protection provided by their organizations' perimeter security systems.

Defending Against Longline Phishing Attacks

When choosing a defense against Longlining and other forms of industrial-scale phishing, it is important to understand the flaws found in existing defense systems that leave them open to circumvention by Longline attacks.

While traditional perimeter security systems are still valid components of an overall security strategy Proofpoint's analysis helps to expose the gaps that such solutions fail to address.

Consider the following security controls and their limitations in detecting and preventing Longline phishing attacks:

- **DKIM and other email signing methodologies:** Unfortunately, while generally excellent at preventing message tampering (we will ignore the short-key DKIM vulnerability here), signing simply validates that the identified sending domain was in fact responsible for the message. But if an email is initiated by a botnet-compromised computer within a valid domain, DKIM and other signing protocols simply validate the attacker's email.
- **Perimeter sandboxes:** While useful for forensic analysis of malware, the efficacy of a perimeter sandbox depends on its ability to recognize and block downloading malware based on its signature, discovered by a prior hapless downloader. This technique won't work for credential attacks, polymorphic malware, or for the 19% of clicks that happen off of the corporate network.

- **URL reputation/classification filters:** While useful for enforcing HR policies and guarding against employee forays into likely dangerous areas of the internet, these legacy web filters are unlikely to be effective against Longlining attacks. As noted by Proofpoint researchers, the attacks compromise known good websites with positive reputations to host their malware. The sites are (correctly) unlikely to be blacklisted based on reputation – and blacklisting at anything less than a full path level would cause challenges.

To more effectively defend against Longline phishing attacks, a protection system would use big data systems to analyze email traffic patterns and characteristics on a per-email-recipent level, and would also use URL redirection combined with cloud-based reputation, sandboxing, and other technologies to unify all knowledge about destination URLS and then (when links are clicked) would allow or deny access on a per-URL basis, at clicktime, every time.

About Proofpoint, Inc.

Proofpoint Inc. (NASDAQ:PFPT) is a leading security-as-a-service provider that focuses on cloud-based solutions for threat protection, compliance, archiving & governance and secure communications. Organizations around the world depend on Proofpoint's expertise, patented technologies and on-demand delivery system to protect against phishing, malware and spam, safeguard privacy, encrypt sensitive information, and archive and govern messages and critical enterprise information. More information is available at **www.proofpoint.com**.

Solutions for Protection Against Phishing and Other Forms of Targeted Attacks

Targeted email attacks represent one of the most dangerous IT threats facing enterprises today. Many of the large, widely-publicized data breaches in recent years have started with a single, carefully crafted email that tricked a recipient to click a link to install malware or surrender their login credentials. Often, these attacks are impossible to detect using conventional reputation, content scanning and sender verification techniques.

Proofpoint Targeted Attack Protection™ takes an entirely new approach, using big data analysis techniques to identify and apply additional security controls to suspicious messages. Targeted Attack Protection represents the industry's first comprehensive, cloud-based solution for combatting targeted email attacks.

To learn more about Proofpoint Targeted Attack Protection, please visit:
www.proofpoint.com/tap

proofpoint™

Proofpoint Limited
200 Brook Drive Green Park Reading,
UK RG2 6UB
Tel +44 (0) 870 803 0704

www.proofpoint.com/uk