proofpoint™

Hook, Line, and Sinker

How Credential Phishing is Changing — How to Stop It

Executive Summary

Phishing is more than 20 years old, but still represents more than 90% of targeted attacks. The reason is simple: it works.

Nearly one in four people who receive a phishing email open it, and more than 10% click on the malicious link or open the weaponized attachment that the phishing email contains. An attacker has to send only 10 messages to have a 90% probability of catching and compromising a user. The average-sized organization loses $3.7 million to phishing scams per year, according to the Ponemon Institute.[1] And those are just the tangible costs.

How phishing works: trends and tactics

While other cyber attacks have become more advanced in a technical sense, phishing has grown more advanced in terms of how it exploits human behavior. They are getting past traditional filters and into corporate inboxes. And they are being clicked, regardless of the amount of user-awareness training. Here are common phishing tactics we have seen in recent months:

- Pretending to be from the targeted users' IT department
- Targeting specific users and departments
- Using weaponized documents embedded with malicious macros
- Working in conjunction with watering-hole attacks

Who attackers target: your people in the cross hairs

Attackers go where the money is. Payment and financial service firms make up 40% of targeted phishing attacks. Internet service firms are also popular targets because compromised domain-name and hosting services enable future attacks.

While attackers send phishing email throughout targeted organizations, click-through rates vary by department and job function. Surprisingly, some the most conscientious workers are among the most likely to click, especially when the phishing email appeals to their sense of urgency, efficiency, and order.

Phishing attacks make money in a variety of ways. Some sell stolen credentials. Others trade in confidential documents. And still others use the stolen credentials to make fraudulent financial transactions.

Why phishing attacks succeed: exploiting human nature

Attackers are using all kinds of tricks to get their weaponized links and documents past organizations' filters and firewalls. They work because they exploit the one vulnerability that can't be patched: your people.

They get past standard email defenses with fast-changing URLs and domains and low-volume campaigns that don't trigger red flags. They impersonate friends and colleagues to lower your people's natural defenses. They use social engineering to create convincing, hard-to-resist emails. And they take advantage of workers' ever-expanding digital footprint to find new ways in.

How to stop phishing attacks: Recommendations

Your people are now the primary exploit target. You need to protect them the way they work and identify assets and risks before you are compromised.

Here are some recommendations for combating phishing attacks:

- Reduce the attack surface. Deploy tools that monitor and analyze messages, URLs, attachments, and user clicks using static code and dynamic behavioral techniques.
- Expand your defense coverage with cloud-based defenses that protect your people wherever they work.
- Take advantage of Big Data and machine learning techniques to predictively catch emerging and never-before-seen attacks before the user clicks.
- Get better visibility into your environment and the broader threat landscape. Real-time threat intelligence and a view of threat activity on your systems help you respond and recover faster. Deploy tools that help you understand who is being targeted by what threats, which threats made it through your defenses, and who has been hit.

With an increasing amount of sensitive and confidential information—and an expanding attack surface of devices, cloud apps, and mobile locations—you cannot afford to rely on traditional defenses. Innovative new tools operate within the work flow, monitoring URLs and attachments, sharing real-time threat intelligence, and watching user activity on and off your corporate network.

Targeted attack protection helps you detect, mitigate, and respond to these threats before they succeed.

1 Ponemon. "The Cost of Phishing & Value of Employee Training." August 2015.

Table of Contents

Introduction

You have just received a Google Drive notification that your colleague shared a document with you. After you log in, you realize that the message was not from your colleague. Unfortunately, you are the victim of a credential phishing scam and it is too late—you have already handed over your password, and your confidential documents may be auctioned or leaked.

In a costlier scenario, you received an email containing details of an invoice or payment issue, or some administrative banking detail. Clicking on the link or attached document conveniently takes you to the bank login screen. Unfortunately, this is a well-crafted website impersonating the bank, and the attackers could now use your credentials to quickly execute multiple funds transfers.

Phishing techniques are more than 20 years old, but over 90% of targeted cyber attacks still use them. At the same time, 61% of breaches begin with credential theft, according to the 2015 Verizon Data Breach Investigations Report.[2] Modern attacks are socially engineered to appeal to those most likely to click, as attackers adjust their lures and learn more about their intended victims.

These highly-targeted attacks are hitting organizations of all sizes. And all too often, they are successful.

The average-sized organization loses $3.7 million to phishing scams per year, according to the Ponemon Institute.[3] Those are just the tangible costs: lost productivity, incident response efforts, cleaning infected systems, and the like. Not counted in these reports is the less measurable—but very real—costs of customer churn, strained business partnerships, and tarnished brands.

This paper describes how today's phishing attacks work, why they are so effective, and how to reduce the chances of one causing lasting harm to your organization.

Phishing Goes Corporate

According to The Human Factor 2016[4] report, phishing campaigns have gone corporate. The most commonly impersonated messages found were corporate-styled voicemail notifications, package delivery updates, invoices, and other financial documents. Campaigns were timed to arrive at their target organization between 9-10 a.m. local time, with the goal of catching people at the start of the workday before IT teams have had a chance to detect and remove malicious messages.

Past Human Factor studies have also shown that phishing operations have shifted targets. Attackers send their malicious emails during the middle of the business day to multitasking staff and middle managers, who are more than twice as likely to click as executives are. People in sales, finance, and procurement are the most likely to click, 50% to 80% more so than other departments.

Why Does Phishing *Still* Work?

Phishing remains a popular attack technique because it works; almost one quarter of all recipients will open phishing messages, and more than 10% will click on the malicious link or open the weaponized attachment. An attacker has to send only 10 messages to have a 90% probability of catching and compromising a user. That's why more than two thirds of all attacks that resulted in a network compromise included at least one phishing scheme.[5] Reports to the Anti-Phishing Working Group—an international consortium of security vendors and law enforcement—more than doubled from 2014 to 2015.[6] This organization found almost 900,000 unique phishing websites last year.

Phishing attacks make money in a variety of ways. Some sell stolen credentials. Others trade in confidential documents. And still others use the stolen credentials to make fraudulent financial transactions.

A recent study from security vendor RSA estimates that businesses worldwide lose more than $400 million per month to phishing.[7] Continued success leads to more phishing activity and innovation; attackers look for new targets, new ones get in on the action, and others evolve their techniques in response to updated defenses.

2 Verizon, "Data Breach Investigation Report", 2015, http://www.verizonenterprise.com/resources/reports/rp_data-breach-investigation-report-2015_en_xg.pdf

3 Ponemon. "The Cost of Phishing & Value of Employee Training." August 2015.

4 Proofpoint, "The Human Factor 2015" https://www.proofpoint.com/us/id/WP-Human-Factor-Report-13

5 Verizon, "Data Breach Investigation Report", 2015, http://www.verizonenterprise.com/resources/reports/rp_data-breach-investigation-report-2015_en_xg.pdf

6 Anti-phishing Working Group "Phishing activity trends report: 1st-3rd quarters 2015" http://docs.apwg.org/reports/apwg_trends_report_q1-q3_2015.pdf

7 http://www.emc.com/emc-plus/rsa-thought-leadership/online-fraud/index.htm

Advanced Attacks, Trends, and Tactics

To be successful, phishing campaigns first have to get through the security filters. Attackers used to rely on the sheer volume of messages to get some through, but they have grown more sophisticated. While large numbers of poorly crafted phishing messages are still sent, the vast majority get filtered out. The few that get through tend to appear more legitimate and trustworthy, and are enabled by social engineering—tricking users into clicking on the malicious content.

In The Human Factor 2016 report, the researchers identify three kinds of victims:

- Enablers are high-volume targets convinced to disable security, click links, or open weaponized files that installs malware on their devices.
- Facilitators are mid-volume targets who are encouraged to input their credentials into very realistic but illegitimate logon screens.
- Gofers are specifically targeted people with appropriate access who are tricked into transfering funds or redirecting shipments by messages impersonate their boss.

Trend: Pretending to be from IT
Looking at a detailed collection of phishing messages that made it through the filters at a major U.S. organization over the past two years, 45% appeared to come from internal IT, and 30% from other parts of the organization. The remaining 25% were distributed among generic subjects, external account verification, urgent issues, financial transactions, and government notices.

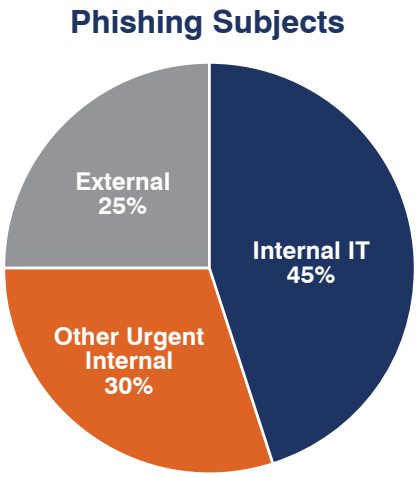

These messages often look familiar to people, covering topics such as mailbox quotas, help desk requests, file sharing, disputed payments or invoices, and received voicemails or faxes. The vast majority of malicious links now lead to credential phishing pages instead of malware. The most effective impersonate popular file-sharing sites.

Trend: Highly Targeted with Social Engineering
A 2015 study looked at over 1,000 phishing messages collected from 2002 to 2014, to determine how they had changed.[8] All of the selected messages made it through the researchers' spam filters. Drawing on prior research, this study analyzed the messages for superficial characteristics and reputation characteristics. Comparing the early group (pre-2007) to the latest group (2014), found minor changes in superficial characteristics, such as grammar and typos. But more significantly, the study revealed big changes in who the messages targeted.

Where earlier messages were widely broadcast, more recent ones are finely targeted to a specific department or person, and they are frequently combined with other social engineering. They use the correct name, for instance, mention relevant projects and co-workers, or even enhance the ruse with phone calls and voicemails. The emails may even appear to come from a colleague or higher-level executive, enhancing their reputation characteristics.

Delivery times are tailored regionally to match the most active part of the workday and workweek. Malicious messages are most likely to be delivered in the morning and most likely to be clicked on Tuesdays.[9] But there was no time of day or day of week that did not have malicious activity.

Social media is also infected with phishing, with attackers creating fraudulent corporate accounts to snag the credentials of victims attempting to contact customer service. More than 55% of Facebook accounts and 25% of Twitter accounts claiming to represent a Fortune 100 company were found to be unauthorized.

Tactic: Macros are Back
Another change is the shift back to malicious attachments, instead of URLs.[10] As people became accustomed to checking URLs in suspicious messages before they clicked, attackers shifted back to attachments. But instead of using obvious executable files, they have resorted to the earlier technique of malicious macros.

8 Yates, D. & Harris, A.L. "Phishing Attacks Over Time: A Longitudinal Study". AMCIS, Association for Information Systems, (2015).
9 Proofpoint, "The Human Factor 2015" https://www.proofpoint.com/us/id/WP-Human-Factor-Report-13
10 http://resources.infosecinstitute.com/spear-phishing-statistics-from-2014-2015/

Most Used Attachment Extensions for Spearphishing in the First Two Months of 2015 (%)

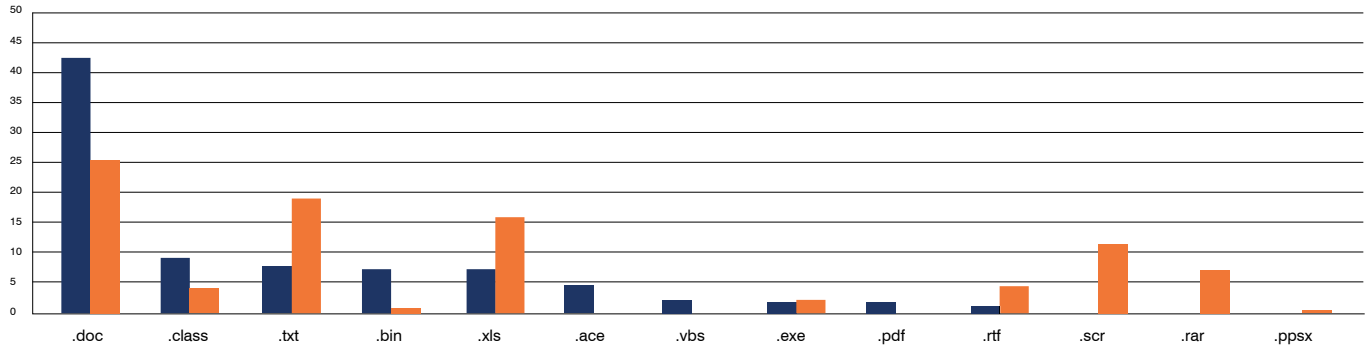

As a security measure, Microsoft changed the defaults to disable macros in Office files more than a decade ago. Users now see a security warning, and macros will not run until the user explicitly allows it. But many businesses use macros to automate and validate Office documents, and users have grown accustomed to ignoring the warning. That complacency, combined with clever social engineering, is leading to more users clicking on documents with malicious macros. Phishing messages with .doc and .xls file attachments have become much more common.

Tactic: Watering Hole Attacks
Most people think of watering-hole attacks as a web-based threat. That's understandable— this kind of attack involves compromising a website frequented by a well-defined target audience to infect visitors with malware. But in many cases, even these involve phishing email.

Once again relying social network research to identify commonly used websites, attackers will send out messages about discounts, coupons, new releases, or other potentially legitimate activities to lure targets to the compromised site.

Who Is Being Targeted?

Unlike other types of cyberattacks, phishing relies more on human behavior than technology exploits.

Over the years, malware writers have adapted to broader and deeper defenses. They have learned to evade sandboxes. They have moved up and down the stack, looking for technological advantage. But while phishing attacks have moved into new message formats and added nicer graphics, they have otherwise remained technologically stable.

That's because susceptibility to phishing is mostly a function of user personality. Even with all of the awareness and user training, people cannot effectively estimate their own risk—and those who underestimate the risk are more likely to be tricked.

Conscientious Workers More Likely to be Caught
Attackers appear to learn from their successes and failures, which would account for some of the trends discussed above.

A recent study found that the most conscientious workers were more likely than their colleagues to click phishing messages that appeal to efficiency, urgency, and order.

Figure X: Example of a phishing email that might trick your best workers

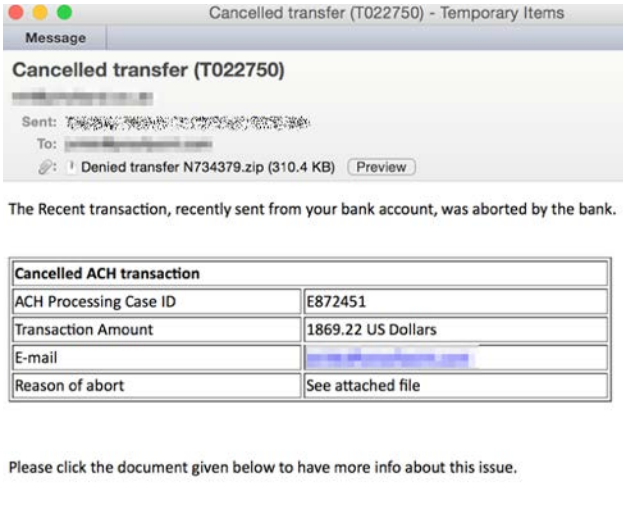

In other words, some of your best employees—those who are hardworking, have the most self-control, and would seem to be the least likely to fall for phishing emails—are really the most likely. Why? Because their personality is being targeted with messages that prey on their work ethic.[11]

Who Gets Phishing Emails—and Who Clicks Them
Phishing messages are being sent throughout the organization, but the volume differs slightly by department. Marketing gets the fewest phishing lures, just under two per day. Sales gets the most, at about three.

Click-through rates are another story. Sales, finance, and supply chain professionals are three times more likely to get phished than those in IT or customer support.

11 Halevi, T., Memon, N., Nov, O., "Spear-phishing in the wild: A real-world study of personality, phishing self-efficacy, and vulnerability to spear-phishing attacks", http://ssrn.com/abstract=2544742, January 2015.

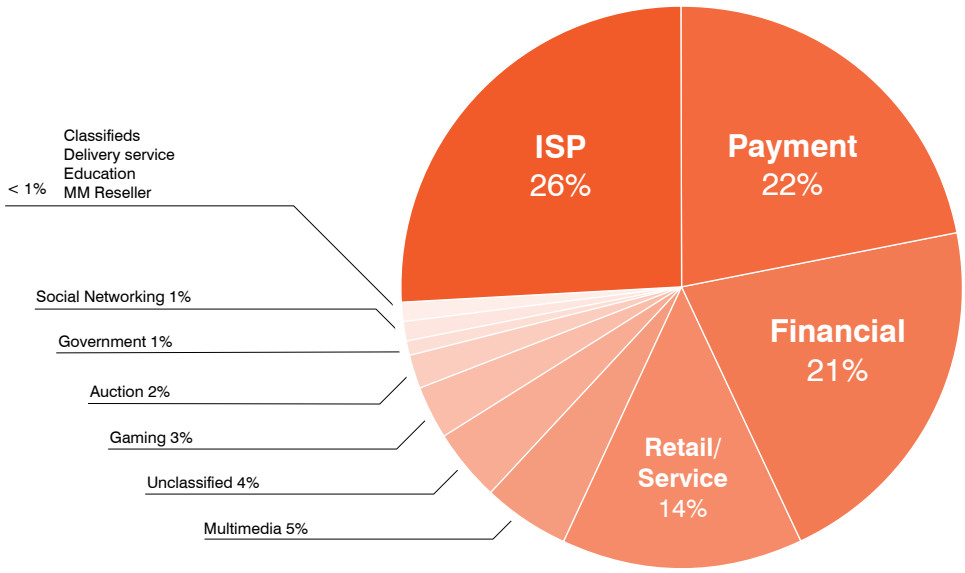

Executives, a prime target in previous years, are now getting no more messages than their managers and staff. But those managers and staff, (in many cases responding to messages that claim to be from the executive) are twice as likely to click as the execs. This most likely reflects the change in tactics towards malicious attachments, macros, and corporate message subjects. In these cases, weaponized documents are crafted to appear as legitimate tasks for the recipient.

Financial Services, Internet Providers, and Retail

Thieves still go where the money is. That's why payment and financial services companies consistently make up over 40% of targeted phishing attacks.[12] Another favorite target is Internet service providers; criminals are looking for account information to send more phishing spam—and if all goes well, get access to domain name and hosting functions to leverage for the next campaign.

Why Attacks Succeed

For phishing to work, it has to first get past the security tools. Attackers are using all kinds of tricks to get their weaponized links and documents past the filters and firewalls that are trying to keep them out.[13] These filters typically look at these things:

- Source of the message
- Keywords in the subject and body
- Reputation scores of embedded URLs
- Malware signatures in any attachments

Getting Past the Filters

In response, attackers are using compromised reputable mail servers and innocuous language to avoid these defenses. Machine learning and crowdsourcing of blacklists have improved filtering of malicious URLs to upwards of 90%. But identifying and disseminating the information can take eight to 12 hours for most tools.

Phishing spammers have responded with fast-changing URLs and domains and much smaller campaigns. More than 15% of global spam is made up of smaller attacks that send only 100 to 5,000 emails. These attacks run for less than 12 hours, and in some cases as short as two hours. So by the time they're identified and blacklisted, the attackers have already ended the attack.

Impersonating Friends and Colleagues

The 10% or less of phishing messages that get through most filters are the best of the best. These messages regularly come from compromised accounts that belong to colleagues, friends, and suppliers—or very good mimics. Some of the most effective message techniques include voicemail and fax attachments, Windows 10 upgrade notices, and invoice or payment issues.

12 Anti-phishing Working Group "Phishing activity trends report: 1st-3rd quarters 2015" http://docs.apwg.org/reports/apwg_trends_report_q1-q3_2015.pdf

13 Alsharnouby, M., Alaca, F., Chiasson, S. "Why phishing still works" International Journal of Human-Computer Studies, 2015.

Impersonating URLs and Websites
While many phishing messages are using malicious attachments, embedded URLs are still being used. Abraham Lincoln once warned that "You can fool all the people some of the time," and indeed, the best phishing sites fool 90% of users.[14] Despite attempts from modern web browsers to provide visual cues of suspicious websites, most people focus only on the address bar.

Spammers are building domains names that are very similar to the legitimate ones. They may use visual substitution techniques, such as substituting "1" for a lower-case "l." They also rearrange letters, and use compound domains like "invoice_legit-firm.com" or "legit-firm.com.phishinglink.com." More than 80% of phishing messages that reach users at this stage have a URL that impersonates a legitimate one with these tricks.

A related technique is graphically emulating the login portion of a commonly visited website. Whether a file sharing site, webmail interface, or financial service, many of these are nearly identical to the original. Busy people will be frustrated by what appears to be a repeated request to login, and quickly enter their credentials. Many of these simply link to the legitimate site, where the user is already logged in, leaving them unaware that they have been phished.

Social Engineering
Supporting all of this offensive trickery is often large-scale social engineering research using public information to refine the attack.[15] They get all the key details right:

- Correctly spelled name and title
- Names of colleagues, suppliers, and customers
- What software and websites the department uses
- Relevant financial information

These details are gathered from social networks, public statements and press releases, search engine keywords, and header details of intercepted messages.

Combined, these bits are assembled into a powerful weapon: familiarity. Once the attackers have built their weaponized messages, they augment them with phone calls, voice mails, mobile messages, or other urgent messages. To the target, they appear to come from an executive assistant, consultant, or other influential person.

Expanding Attack Surface
Finally, attackers are taking advantage of the ever-expanding attack surface that we are all creating. Multiple devices, numerous work locations, and collaborative web apps all provide additional openings. Smaller screens make evaluating the legitimacy of potential phishing messages or websites harder. And the constant pressure to respond makes it more likely that someone will act on a seemingly friendly message will be acted upon.

Attackers also want to catch people when they are away from their core network protection, getting through the much weaker defenses—if they exist at all—at home or on the road.

Phishing education and awareness are important. But you cannot educate your way to 100% coverage. With the average time from click to compromise less than one minute, you cannot afford to rely on legacy endpoint and firewall defenses that wait until after the click. You need to take action long before the phishing messages are delivered.

What you need to defend against today's phishing attacks

Modern phishing attacks are getting through traditional defenses and human detection. That's why today's advanced attacks call for new, innovative solutions. These tools must reduce your attack surface to minimize the threat, use cloud-based intelligence and analytics to operate at the necessary scale, and accelerate the response capabilities so that you can prioritize your efforts and recover quickly.

Reduce the Attack Surface
The first step in defending against modern phishing attacks is reducing your attack surface. To do this effectively, you need to know what your users are doing, their normal behavior patterns, and when sensitive or confidential information might be compromised.

Social engineering and human exploits are an increasing part of successful attacks, whether delivered via email, mobile apps, or social networks. That means your solution must continuously monitor and evaluate messages, URLs, attachments, and user clicks.

Attackers' dynamic code, meanwhile, has greatly diminished the benefits of signature analysis. Today's defenses must check multiple characteristics. They must be able to quickly detect known and emerging malicious addresses. They must analyze code snippets for hidden malicious behavior. And they must dynamically observe actual behavior in a sandbox—one that can simulate a real user to defeat sandbox-evasion techniques such as sleep timers and mouse movement.

14 Dhamija, R., Tygar, J.D., Hearst, M. "Why phishing works" Proceedings of CHI-2006, Conference on Human Factors in Computing Systems, 2006
15 Infosec Institute, "Phishing attacks using public data", 2016, http://resources.infosecinstitute.com/phishing-attacks-using-public-data/.

Cloud-Based for Scale and Scope
The next step is expanding your defense coverage. Cloud sharing and collaboration tools, social networks, bring-your-own-device, and mobile work locations are taking more and more work off the corporate network—it's no wonder one in five clicks on malicious links occur outside of your network perimeter.

Phishing attackers see great opportunities away from corporate eyes, as smaller screens and greater time pressure improve the click rate. Users and attackers have learned how to get around proxy servers. That means you need to dynamically rewrite links in email messages and test them every time they are clicked. Mobile apps often permit risky behavior, leading to leaked data or stolen credentials.

Cloud-based solutions enable you to quickly deploy and scale to cover tens or hundreds of thousands of users, regardless of their location, device, or cloud application. As more and more applications become cloud-based or mobile-enabled, defenses need to move the same way.

Predictive Defenses
Many organizations are using Big Data techniques to better understand customer behavior. Cloud-based threat intelligence and statistical modeling deliver similar benefits against advanced attacks.

Machine learning can build statistical models of each user, flagging messages that are anomalous or suspicious. Scoring these in real time against continuously updated threat models, existing phishing campaigns, content inspection, and other characteristics can predict malicious messages and URLs. The result is a suite of defenses that can catch emerging or never-before-seen attacks before the user clicks on a malicious link or opens a weaponized attachment.

Superior Intelligence and Visibility
Unfortunately, as long as humans can be exploited, no security system will be 100% effective. So you also need to be able to see ongoing attack campaigns and what threats are being directed at your people. Web-based and graphical dashboards that show you:

- Who is being targeted by what threats
- When the attackers entered your environment
- Which ones made it through your defenses
- Who has been hit

Armed with this accurate and detailed information you will be able to accelerate your response, prioritize incidents, and recover quickly.

The Benefits of an Effective, Integrated Defense

Throughout your organization, you have an assortment of high-value people, information, and digital assets. Modern phishing attacks are getting past traditional defenses. A modern, complete defense suite can protect your sensitive and confidential data from targeted attacks.

Cloud-based defenses deploy quickly, expand to handle the breadth and depth of your applications, and deliver the Big Data processing capacity, with less cost than on-premises solutions. And integrated incident response dashboards provide faster insight, reducing your reaction time and minimizing exposure.

How Proofpoint Can Help

Proofpoint provides a cloud-based suite of cybersecurity products that help organizations efficiently detect, mitigate, and respond to advanced threats by protecting in today's workflow: email, social media, and mobile applications. Proofpoint Targeted Attack Protection (TAP) in Email provides unparalleled defenses against polymorphic malware, weaponized documents, and credential-stealing threats.

Protect in-flow of how people work today

To protect people from advanced threats, IT needs to align with how people work today.

That means:

- Analyzing the data associated with business-critical tools that drive your operations
- Enabling you to understand your risk and mitigate your attack surface.
- Protecting your people on any device, on and off the corporate network

Deploy quickly in the cloud; realize value immediately

The cloud architecture allows you to deploy quickly and derive value immediately.

With TAP, you can:

- Receive instant software updates for rapid incorporation of new defenses
- Rapidly deploy to hundreds of thousands of users in days
- Adopt the only viable solution for advanced threat protection in key cloud technologies

Detect known and unknown threats using sophisticated, adaptable techniques

Organizations need security solutions that stay ahead of the changing threat landscape.

That means:

- Analyzing threats in multiple stages using multiple techniques that cover the entire attack chain
- Mitigating and responding effectively with a broad understanding of the attack surface
- Continuously adapting techniques to stay ahead of attackers

Apply superior security intelligence for faster detection and response

Proofpoint provides unmatched insight derived from community-based intelligence.

This means you can:

- See attack campaigns affecting entities around the globe across diverse industries
- Connect the dots using a threat graph with more than 300 billion data points, learning more about attacker tradecraft and making the next attack even easier to catch
- Get the forensic information you need to respond to threat alerts effectively
- Benefit from the integration of Proofpoint Emerging Threats Intelligence data to enrich your threat insight

Conclusion and Recommendations

Phishing has changed. Your people are now the primary exploit target. You need to protect them the way they work and identify assets and risks before you are compromised.

Today's cyber attacks require a solution that you can deploy quickly and put into action in days, not months. You need the capacity and scale of cloud-based analytics to accurately make threat associations across a massive real-time threat intelligence database. And you need the tools to predictively identify unknown threats and quickly learn from each attack.

Attackers are getting better at impersonating your people for their targeted attacks. Targeted attack protection helps you detect, mitigate, and respond to these threats before they succeed.

To learn more about how Proofpoint can help your organization detect and stop credential phishing. Contact us today for a free assessment offer at www.proofpoint.com/cybersecurity-assessment.

about proofpoint

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.

proofpoint™

892 Ross Drive
Sunnyvale, CA 94089

1.408.517.4710
www.proofpoint.com