Email Security for Office 365

It's broken. Here's how to fix it.

Data breaches are increasing.
Phishing is on the rise, costing businesses $1.6 million per attack.
Spear phishing is a top threat, with losses of $1.2 billion in 2018.
Malware is constantly mutating and proliferating.

What's the common thread behind all of these trends? Email.

VadeSecure
Predictive Email Defense

Contents

Introduction 1

Phishing 3

Spear Phishing 5

Business Impacts of Email Attacks 8

Why Are So Many Organizations Vulnerable to Email-Borne Threats? 10

The Pros and Cons of Microsoft Exchange Online Protection (EOP) 11

Going Beyond Signature-Based Protection 12

    Vade Secure's Solution 13

Conclusion 15

INTRODUCTION

The perimeters of vigilant organizations are reasonably tight. Firewalls are in place, servers are patched, and physical security is in place. However, email is a gaping hole in your network defenses.

Email is the vector for virtually all the bad things that keep you up at night.

What's worse is that cybercriminals are using email to target the weakest link in an organization's cybersecurity chain: humans.

If you think your organization is safe from email-borne attacks because you have enabled Office 365 email security add-ons like Exchange Online Protection (EOP) or Advanced Threat Protection (ATP), think again. While these security tools are effective at blocking massive spam waves and known threats, they will not reliably stop highly targeted phishing, spear phishing, or zero-day attacks.

According to the 2018 Data Breach Investigations Report, phishing and pretexting represented 93 percent of social attacks in 2018.

Email security is clearly broken.

The problem is twofold:

1. **Technology:** Most email "security" systems are just glorified spam filters. They were designed to stop known mass email attacks. The underlying architecture of these solutions isn't suitable to catch zero-day threats or one-off spear phishing emails.

2. **People:** Many employees will click on or respond to a well-crafted phishing or spear phishing email if it lands in their inbox. Only 17 percent of phishing campaigns are reported by users, and most people who are phished do not report the campaigns. When a phishing campaign is launched, IT has approximately about two minutes before a user clicks on the email. [1]

We'll walk through the scope of the problem and then discuss how you can quickly close existing security holes in your Office 365 environment.

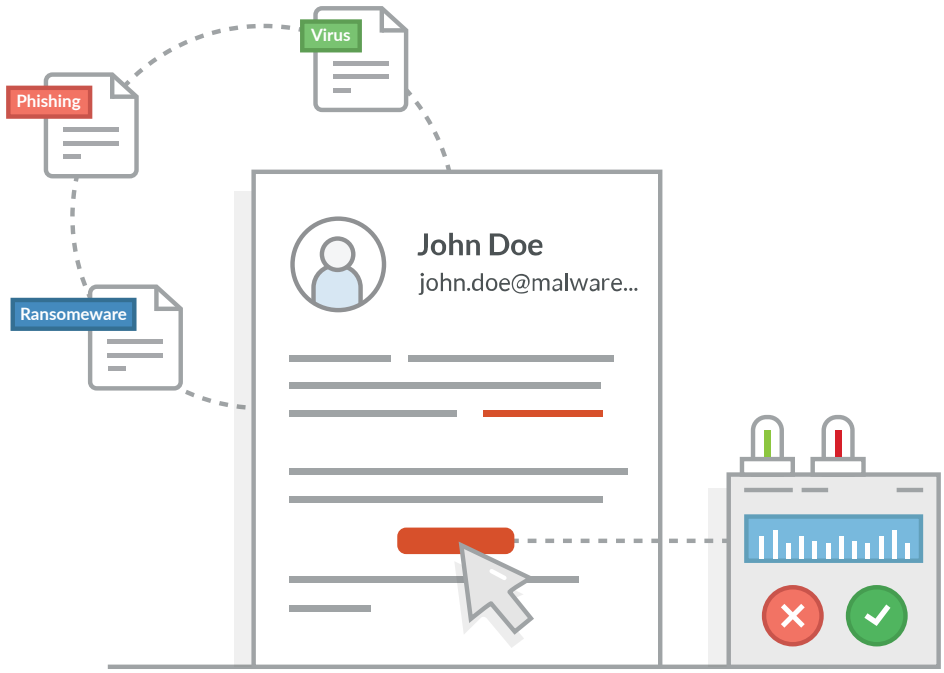

[1] Aberdeen Research. "Reduce the Risk of Phishing Attacks: The Race is On." December 2018.

PHISHING

Phishing is a hacking technique that “fishes” for victims by sending them deceptive emails. (The “ph” replaced the “f” in homage to the first hackers, the “phone phreaks” of the 1960s and ’70s.) Virtually everyone on the Internet has seen a phishing attack. Phishing attacks are emails that request confidential information or credentials under false pretenses, link to malicious web sites, or include malware as an attachment. In the past, phishing emails were sent to hundreds—even thousands—of recipients in mass waves. Today, phishing attacks are sent in lower volumes—typically to three or fewer employees—with sophisticated levels of personalization, and with more precision.

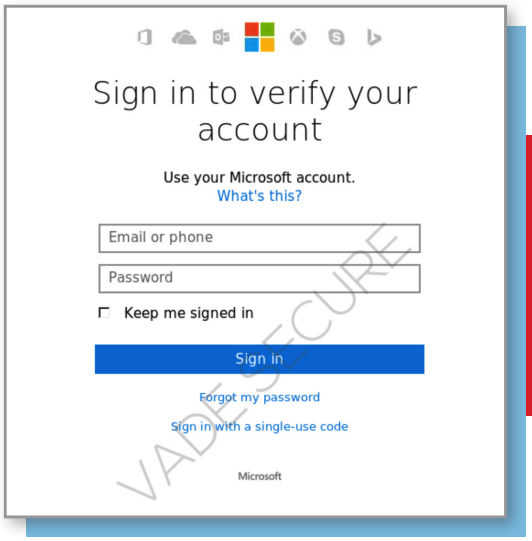

Figure 1 - Busy people **will** fall for realistic login screens ... and once they enter their credentials, your network is toast.

Many phishing sites look identical to the legitimate brand or site they are impersonating. While banks and online payment vendors like PayPal have always been a top target for spoofing, Microsoft has been the #1 impersonated brand for five quarters—thanks to the lucrativeness of Office 365 applications, including SharePoint and OneDrive. With access to a corporate Office 365 account, cybercriminals can conduct spear phishing campaigns from the inside, including requesting bogus pill payments or direct deposits changes, or even selling the company’s global address list to spammers. Microsoft phishing pages are highly sophisticated and extremely difficult for the average user to identify as phishing. Often, the only difference in Microsoft phishing pages is a slight (and easily missed) difference in the URLs.

Victims can easily be duped into disclosing credentials or confidential information to the hacker if they can be induced to click the link in the email. Common Office 365 phishing techniques include emails asking users to log in to their Office 365 accounts to update account information, retrieve a shared file, or re-enter their password to regain access to the platform. Once a user has entered their password on the phishing page, their credentials are harvested.

Blacklisted phishing sites can often get by standard filters through the technique of time-bombing the URLs. The hackers will include a clean URL in the email to get past the filters and then redirect to a malicious site once the email has been successfully delivered.

Although malware is harder to get past filters, recently discovered and zero-day malware stands an excellent chance of getting through standard filters (and being clicked on), especially if the malware is hidden in a non-executable file like a PDF or Office document. This is how many of the recent ransomware attacks were propagated.

The share of inbound phishing emails in Office 365 increased 250 percent between January and December 2018. [2]

As we will see, all of these attacks become much more devastating when carefully customized and individually sent via a spear phishing email.

[2] Microsoft. "Microsoft Security Intelligence Report. Volume 24. January – December 2018."

SPEAR PHISHING

Spear phishing is an enhanced version of phishing that takes aim at specific employees of a targeted organization. The motivation is usually financial, with the most common attacks coming in the form of wire transfers or gift card requests, or requests to divert direct deposits or vendor payments into fraudulent accounts. In contrast to the mass email approach used in phishing, which might see hundreds of thousands of messages sent to random recipients within the space of a few hours, spear phishing is methodical and focused on a single recipient. Often the initial email, known as "pretexting," will contain no URL or attachment. Rather, it will simply try to provoke a response and develop a "conversation" to lull the recipient into thinking the sender is legitimately whomever they are posing as. Only later will the hackers request confidential credentials or information, or payment via wire transfer or gift cards.

The additional customization and targeting of a spear phishing email, along with the lack of easily recognized blacklisted URLs or malware, will generally bypass standard email filters.

To show how the spear phishing process works, let's explore an attack on a hypothetical widget company called Widget Co., which has 500 employees in different cities. Hackers are interested in getting access to Widget Co's database of employee records. They can harvest the employees' confidential information, such as Social Security numbers and direct-deposit bank accounts, and sell them on the black market to identity thieves.

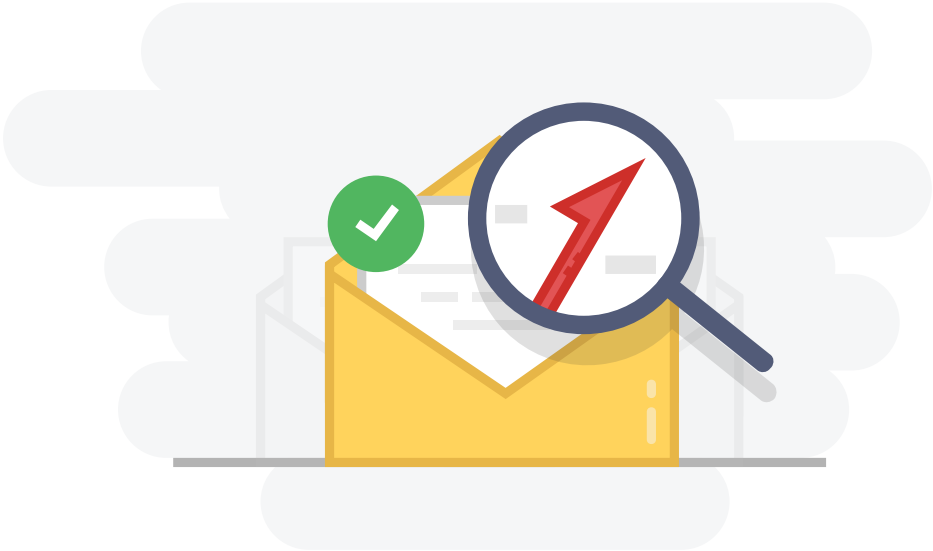

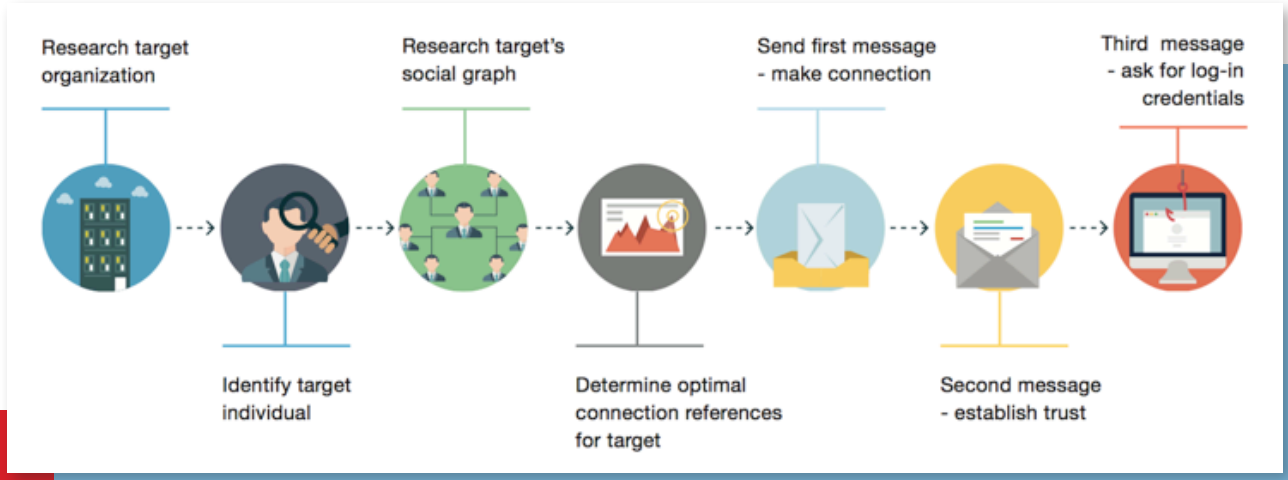

Figure 2 - The progression of a spear phishing attack, starting with research of the target organization and identification of a specific individual inside the organization, followed by a series of emails intended to build trust with the target.

Figure 2 shows a typical progression of a spear phishing attack. The attacker's first step is to research Widget Co. to get a sense of how they can best mount a successful spear phishing attack. After cataloguing the executives in the "Our Team" section of the Widget Co. website, the attackers create a cross-reference of social graphs, using Facebook and LinkedIn accounts to build lists of who knows whom inside Widget Co. Then, by piecing together the social information, the attackers are ready to go spear phishing.

The attackers find an HR employee at Widget Co. named John Smith. Posing as Mr. Smith, the hackers target Smith's Facebook friend and colleague, Jeff Jones, an HR manager at Widget Co. To build trust in the faked email address, the hacker posing as Mr. Smith sends his "friend," Mr. Jones, a note asking about the family vacation he is currently on (according to pictures posted to Facebook). If Mr. Jones responds, the hacker is off to a good start. He's successfully impersonating another Widget Co. employee and is starting to build trust in the faked email with his target. Mr. Jones replies and says he is enjoying his time away with his family. The two continue to banter about Mr. Jones' family vacation as well as things going on in the office, including the names people that have been researched and associated with the social circle.

How can the attacker get away with this? Doesn't Mr. Smith have a unique, domain-specific email through Widget Co.? Yes, he does. However, due to Widget Co.'s "Bring Your Own Device" (BYOD) policy, employees are able to use personal mobile devices to send messages to one another. In this case, the attacker knows from LinkedIn that Mr. Smith's personal email address is johnsmith1@gmail.com. The attacker creates a Gmail account for johnsmith.1@gmail.com. Mr. Jones doesn't notice the difference, and the stage is set for the real attack.

The hackers know from LinkedIn that Jane Doe is a new employee working with Mr. Jones. The hacker posing as Mr. Smith sends to Mr. Jones a pdf file of "new employee paperwork" that actually contains key logging malware. If Mr. Jones opens the file, his device is instantly infected, his credentials vacuumed up, and the network is breached.

Alternatively, the fake Mr. Smith could send a note that says, "Hey, Jeff — I'm on the golf course, but I need to call the bank and make sure Jane Doe's retirement plan is all set up. I can't remember the login for the employee database system — can you help me out?" If Mr. Jones shares his login for the database, the hacker is inside. Either way, the phisher can collect Mr. Smith's login credentials — a free pass to invade the Widget Co.'s private networks. Any confidential employee data is at risk of being improperly accessed.

In this case, we used an HR example, but it could just as easily have been in corporate finance, marketing and sales, IT, or any other department. Most employees have more than enough personal information about them in the public realm to allow their identity to be utilized to swindle another employee and compromise your network.

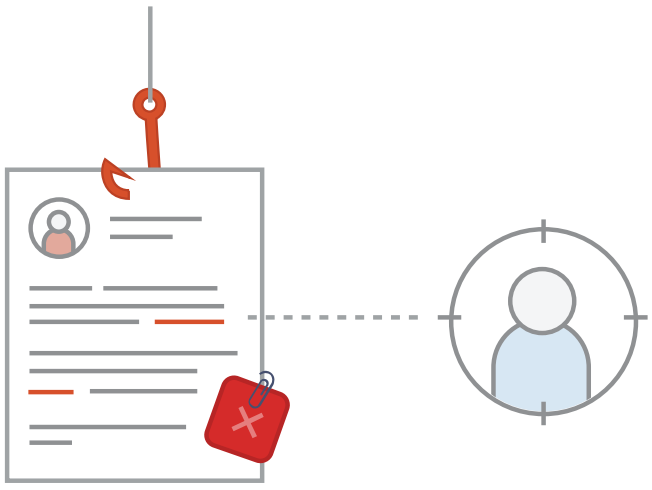

BUSINESS IMPACTS OF EMAIL ATTACKS

The impact of such attacks can vary but will generally increase with the sophistication of the attacker and the size of the target. On average, a successful phishing attack costs a mid-sized business $1.6 million.

Consider the financial repercussions of a hacker gaining access to your critical data. What can he or she do with it?

Spear phishing attacks are often just the first part of a multiphase attack. Once inside, the hackers can do devastating damage by accessing confidential customer lists, intellectual property, and emails, and even deleting critical data or encrypting it with ransomware.

Companies that fall prey to hacking enabled by spear phishing face risks of reputation damage, loss of market value, competitive disadvantage, legal liability, and compliance problems. With GDPR affecting not only businesses in the EU, but also any businesses that collects data on subjects in the EU, a data breach comes with the risk of financial penalties. Businesses can be fined up to 4 percent of global turnover for not contacting authorities within 72 hours or not conducting an impact assessment after a breach. Additionally, individual executive careers can suffer in the wake of such events.

Example: Equifax

In 2019, Moody's, a credit rating corporation, downgraded Equifax from stable to negative as a result of the 2017 Equifax data breach that exposed the personal data of 140 million Americans. This is the first time a cyberbreach has been factored into a business's overall credit rating, but experts say it will become increasingly common. Not only did Equifax take a hit to their credit rating, but their reputation was—and remains—deeply injured as a result of the breach. Additionally, the company was hit with a $690 million charge for the breach. [3]

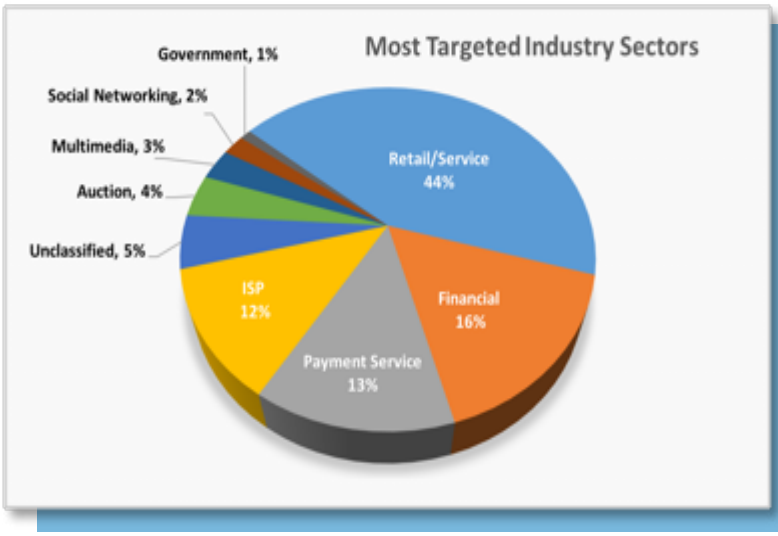

Figure 3 – Most targeted industry sectors, Q1 2019

(Source: APWG Phishing Activity Trends Report, Q1 2019)

[3] Fazzini, Kate. "Equifax just became the first company to have its outlook downgraded for a cyber attack." CNBC.com, May 22, 2019.

Risks by Industry

Education: Local school districts as well as local and state colleges and universities have seen an alarming increase in email attacks, with one school district in Atlanta, GA, reporting that they receive no less than 3,000 attacks per day.[4] Storing the personal records of not only staff but also students, K-12 and secondary education institutions working with small budgets are highly vulnerable to email attacks targeting employee and student data. In a 2019 test of university cyber defenses in the UK, ethical hackers managed a 100 percent rate when attempting to hack university systems—spear phishing was named as a primary tactic in the hack. [5]

Financial Services: Financial firms must manage spear phishing risks that can result in theft of insider trading information, personally identifiable information, credit card numbers, bank account information, and more. The impacts include financial loss, legal liability, and regulatory penalties. In 2017, cybercriminal group London Blue conducted mass spear phishing attacks against financial executives in the US, UK, Spain, the Netherlands, and New Mexico. The group managed to steal sensitive data and funds through wire transfers. [6]

Local Government: With fewer resources than large, private organizations and operating with tight budgets, local governments known for having poor defenses are top targets for email-borne cyberattacks. Malware and ransomware are especially dangerous to local governments. Several recent, high-profile attacks in Baltimore, MD, and Greenville, NC, have debilitated 911 systems and locked government administrators out of city computer systems and servers for weeks and even months. [7]

Retail: As several large-scale hacks have shown, retailers are vulnerable to attacks that leak customer data— including credit-card holder information. This puts them at odds with PCI regulations, which carry fines and costly compliance remediation penalties. They also risk loss of consumer trust and brand value. Retailers also face an indirect risk from spear phishing, which is liability for fraudulent sales made with stolen credit card numbers. [8]

Intellectual Property-Based Businesses: For businesses such as pharmaceuticals and technology, where digital information may represent massive investments, spear phishing can have an especially costly impact. Competitors can gain access to confidential intellectual property that took years and cost billions of dollars to develop.

Manufacturing and Defense: Strategic manufacturing industries and defense contractors are vulnerable to corporate espionage, both private and sovereign. Defense companies are frequent targets of sovereign attackers, such as the cyber warfare units of foreign powers. A serious incident could endanger national security and affect a company's ability to secure further defense contracts.

Health Care: HIPAA-regulated entities are bound by extensive, rigid compliance guidelines and face stiff financial and legal penalties for data breaches. There are reputation risks given the sensitive nature of leaked personal health information. In 2017, the French insurance fund, Amelie, was targeted in several high-profile phishing attacks that tricked users into entering account credentials to either collect reimbursements or enter personally identifiable information. [9]

[4] McCray, Vanessa. "Cyberattacks Increasingly Target Student Data." Government Technology. December 28, 2017.
[5] Coughlan, Sean. "Hackers beat university cyber-defences in two hours." BBC News. April 4, 2019.
[6] Jay, Jay. "London Blue evolves its tactics from phishing attacks to impersonation fraud." SC Magazine UK. April 5, 2019.
[7] Ian Duncan and Christine Zhang. "Analysis of ransomware used in Baltimore attacks indicates hackers needed 'unfettered access' to city computers." Baltimore Sun. May 24, 2019.
[8] APWG. "Global Phishing Report 2Q 2016."
[9] Dumons, Olivier. "A refundable scam mimics the messages from Amelie.fr." May 23, 2017.

WHY ARE SO MANY ORGANIZATIONS VULNERABLE TO EMAIL-BORNE THREATS?

The problem is that native Office 365 filtering systems, such as EOP, will not catch targeted phishing and spear phishing emails. The architectures of these email security systems (as well as the vast majority of other standard email security systems) were originally built to fight spam. Therefore, they focus on mass emails, using reputation and signature techniques to block suspicious emails and *known* malware attachments and phishing URLs.

These processes, while highly successful in fighting spam, are not very useful in the struggle against dynamic phishing and spear phishing. Phishing emails sent at low volumes and that feature sophisticated obfuscation techniques, including URL redirects, will bypass a reputation and signature-based filter that is scanning for known phishing links. In some cases, phishing links are not in the body of the email but in an attachment or a zipfile, which would not be detected by sandboxing. Additionally, in short-wave attacks, phishers often change their mode of attack every few emails, and a one-off, well-written spear phishing email that doesn't include links will generally get past most corporate spam filters.

Standard email security is OK at blocking mass spam attacks ...

Spam-derived email security actually works *OK* for most mass-emailed, *known* phishing attacks...

However, signature-based email security is completely ineffective against sophisticated, highly targeted phishing and one-off, targeted spear phishing attacks and zero-day malware—the primary threats today to your network security.

Today's businesses need a purpose-built email *security* system that will stop all types of email-borne threats—not just a glorified spam filter.

THE PROS AND CONS OF MICROSOFT EXCHANGE ONLINE PROTECTION (EOP)

As stated, EOP can be moderately effective against *known* threats. The problem is that it is not effective at fighting unknown threats … whether from zero-day code buried in an Excel file or a business email compromise (BEC). In fact, according to a 2018 SE Labs report, EOP had a protection accuracy rating of -15 percent and a total accuracy rating of only 8 percent.[10]

Here's what EOP is missing from a security perspective: [11]

- The ability to identify new and evolving threats for which it doesn't have a known signature.
- Real-time URL and page exploration to ensure links are safe and guard against time-bombed URLs.
- Real-time behavioral analysis to detect zero-day malware attacks—without the latency issues caused by sandboxing.
- The ability to scan emails sent within the organization.
- Robust impersonation detection and natural language processing to detect one-off spear phishing emails

Many other email "security" systems are based on spam filtering technology and have the same security flaws in that they can't reliably identify unknown threats, such as spear phishing attacks or unknown malware posing as a non-executable file. Although some of these vendors claim to have some basic analysis that can detect business email compromise, they are only able to detect clumsy fraud, like when there's a difference between the from and reply-to domains or an internal domain. These types of analysis are very basic and can easily be circumnavigated by even moderately sophisticated hackers.

On its own, EOP is not enough to protect from advanced phishing, spear phishing, and malware in Office 365. To achieve full protection requires augmenting EOP, rather than displacing it, as is the case when using a Secure Email Gateway. IDC recommends a layered approach that complements and layers effectively with EOP. "Reputation-based defenses are effective against known threats," says IDC analyst, Konstantin Rychkov, "but the growing sophistication of messaging attacks requires layered controls for unknown, highly dynamic threats and/or BEC impersonation attacks." [12]

[10] SE Labs. "Email Security Services Protection. December 2018."
[11] EOP is also missing some nice-to-haves such as effective graymail (low-priority email) classification, one-step unsubscribes, and archival capabilities.
[12] IDC Analyst Connection. Sponsored by Vade Secure. "Email Security: Maintaining a High Bar When Moving to Office 365." EMEA44752219. January 2019.

GOING BEYOND SIGNATURE-BASED PROTECTION

Vade Secure recognized several years ago that the standard signature and reputation-based security tools were insufficient to protect organizations from a dynamic, rapidly evolving threat landscape in Office 365. What is required is predictive email defense that can recognize brand new threats based on previous patterns. In short, we needed artificial intelligence that has been trained specifically to find these zero-day threats.

Vade Secure protects more than 600 million mailboxes worldwide and processes billions of emails every day. Our 10+ year relationships with the world's largest ISPs provides Vade with unparalleled access to data and helped us develop industrial-scale solutions built for speed and performance.

This gave us a very large data set to start training artificial intelligence and machine learning algorithms to identify malicious emails, phishing pages, and malware in Office 365 environments, so that they can be blocked from the very first email. Our system is capable of reliably identifying one-off spear phishing emails, sensitive data requests, and zero-day malware hidden in executable files, PDFs, Office documents, and more.

Vade Secure's machine learning predictive models are being constantly fine-tuned to ensure a high degree of accuracy. New rules and information are constantly being fed into the system by our 24/7, follow-the-sun security operations centers.

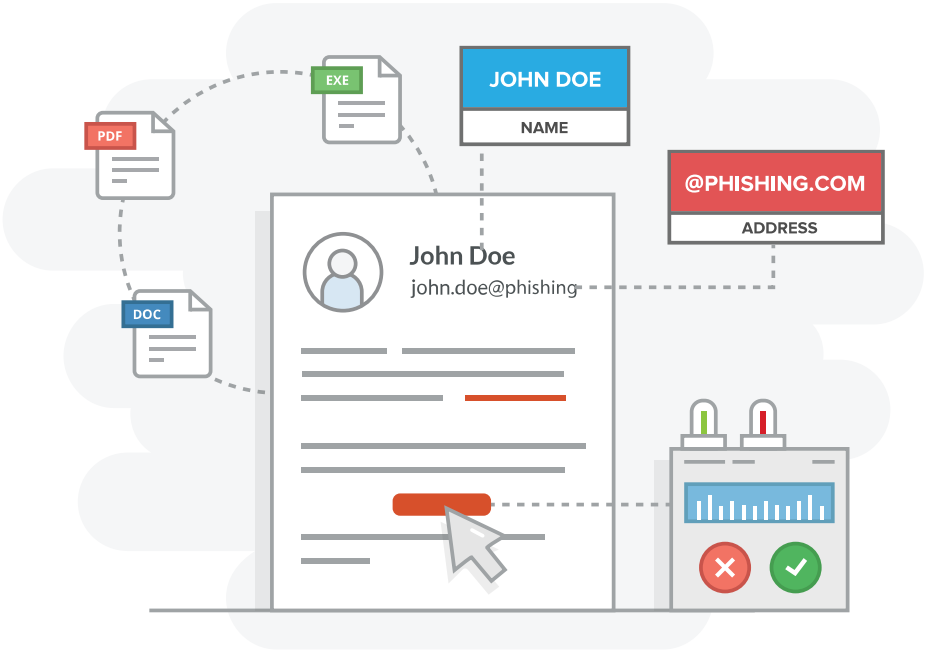

Vade Secure's Solution

Vade Secure for Office 365 provides the most robust email security solution on the market, including protection against phishing, spear phishing, and malware, as well as spam control and graymail classification. Natively integrated with Office 365, it layers with and complements rather than displaces EOP.

- **Filtering:** Using its predictive engine, Vade Secure scans for anomalies, inconsistencies, and malicious behavior in the message's structure, content, attachment, and links.
- **Time-of-Click Anti-Phishing:** Using smart patterns and machine learning algorithms, Vade Secure crawls the URL and webpage, following any redirections to reach the final page and determine whether it's fraudulent. Unlike most URL exploration engines, we explore the URL both when it first enters the system and any time a user attempts to click on a link, thus defeating time-bombed URLs.
- **Auto and Manual Remediation:** Vade Secure augments threat detection with post-delivery threat remediation. With a real-time view of global threats, the engine is continuously learning and will automatically remove threats from user inboxes. Admins can also manually remediate messages with one click.
- **Banner-based Anti-Spear Phishing:** Vade Secure builds a technical profile for each individual with which your employees communicate. Our Identity Match™ system considers hundreds of subtle technical and behavioral factors to determine if the sender is who they claim to be to protect against email imposters. Natural language processing is used to detect malicious patterns, such as flag words of phrases, while unsupervised anomaly detection looks for anomalies such as senders who do not match the organization's entity model. Upon detecting any anomalies, the solution displays a banner within the email alerting the user that the message might be malicious.
- **Behavioral-based Anti-Malware:** Going beyond simply scanning email attachments, Vade Secure performs a comprehensive, 360-degree analysis of the origin, content, and context of incoming emails and their attachments. Supervised machine learning algorithms holistically analyze more than 47 features of the email, attachment filenames, and their content to identify and block both known and unknown malware and ransomware.
- **Human Intelligence:** Vade Secure mans a 24/7 global threat intelligence center with email security experts. They constantly monitor the information that comes in so that we can identify and block new and emerging threats.

- **Native API Integration:** Because of its native integration with Office 365, Vade Secure for Office 365 sits inside the Microsoft tenant, requiring no MX changes and making it invisible to hackers. Easy to deploy and set up, it layers with EOP, provides insider attack protection, and requires no external quarantine.
- **Spam Control:** Vade Secure achieves a 99.99 percent catch rate with a false-positive rate of only 0.0625 percent.
- **Graymail Management:** Vade Secure automatically classifies low-priority messages (e.g. newsletters, promotions, social notifications), while one-click safe unsubscribe easily eliminates unwanted communications, allowing users to have a cleaner inbox.

CONCLUSION

Defending against phishing, especially the spear phishing variant, is a never-ending process. Each day brings fresh versions of the threat to employee inboxes at every organization. Countermeasures must be strong but also adaptable. Artificial intelligence and specialized email security are critical to keeping your organization safe.

Just one click is all it takes to cause significant financial and reputational damage to your business and your clients.