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 2017 losses of $675 million.
Malware is constantly mutating and proliferating.

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

VadeSecure
Predictive Email Defense

Contents

INTRODUCTION

93 percent of all network breaches include a phishing or spear-phishing attack.

For the most part, 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 that 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), you need to 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 State of the Phish™ Report 2018, 76% of organizations experienced phishing attacks in 2017. 97 percent of these organizations already have an email/spam filter in place.

Email security is clearly broken.

The problem is twofold:

1. **Technology:** Most email "security" systems are really 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. Despite education efforts, 20 to 30 percent of recipients open standard phishing messages and 12 to 20 percent of those click on any enclosed phishing links. These already high rates more than double when looking at spear-phishing emails. [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.

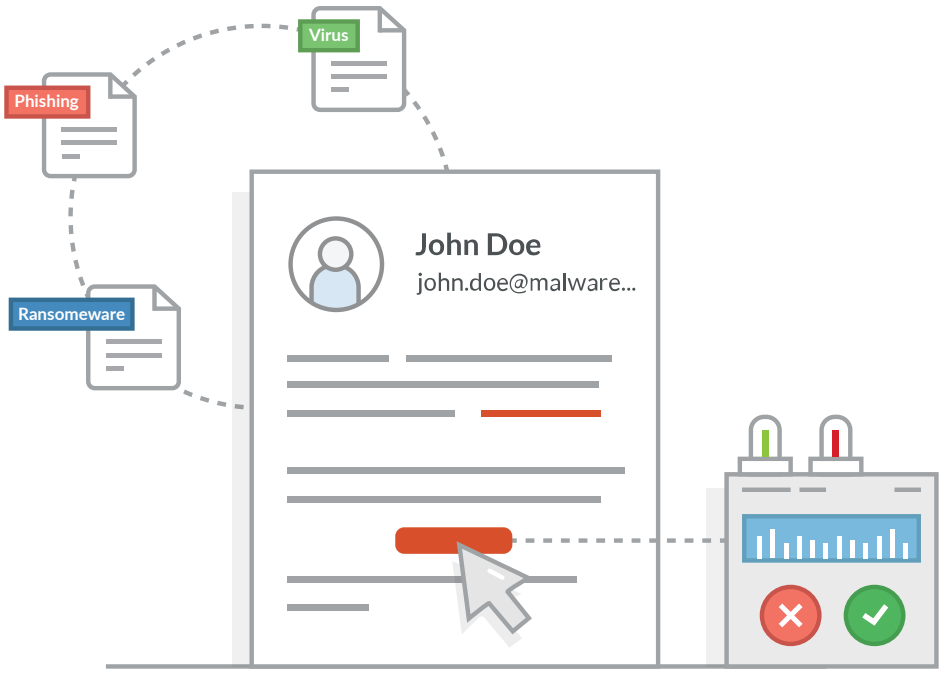

1. The low figures are taken from Verizon's 2016 Data Breaches Report, and the higher numbers are from an August 2016 study by Friedrich-Alexander University, but many other studies show similar results.

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 mass emails that request confidential information or credentials under false pretenses, link to malicious web sites, or include malware as an attachment.

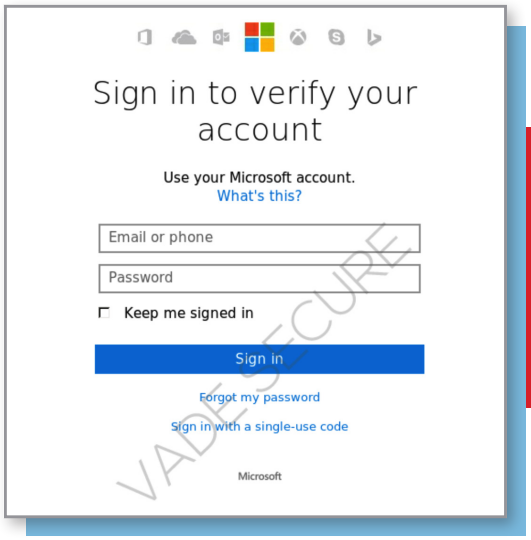

Figure 1 - Busy people **will** fall for realistic login screens ... and once they enter their credentials, your network is toast.

Many phishing sites look identical to the legitimate brand or site they are impersonating. Often, the only difference in many spoofed websites is a slight (and easily missed) difference in the URLs. Visitors can easily be duped into disclosing credentials or confidential information to the hacker if they can be induced to click the link. Even blacklisted phishing sites can often get by standard filters through the technique of time-bombing the URLs. The URL will lead to an innocent URL initially to get past the filters and then redirect to a malicious site.

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.

Despite the lack of personalization, an astonishing 20 percent of recipients will click on basically anything that makes it to their inbox.[2]

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

2. Ibid.

SPEAR PHISHING

Spear phishing is an enhanced version of phishing that takes aim at specific employees of a targeted organization. The goal is usually to gain unauthorized access to networks, data, and applications. In contrast to the mass email approach of 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 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 send a booby-trapped URL or attachment.

The additional customization and targeting of a spear-phishing email, along with the lack of easily recognized blacklisted URLs or malware, will generally get it past standard email filters. What's worse, this same customization results in click-rates in excess of 50 percent![3]

To show how the spear-phishing process works, let's explore an attack on a hypothetical widget company called Widget Co., which has 10,000 employees spread out over five campuses in different cities. The company employs more than 500 administrative employees. Hackers are interested in getting access to Widget Co's database of hundreds of thousands 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.

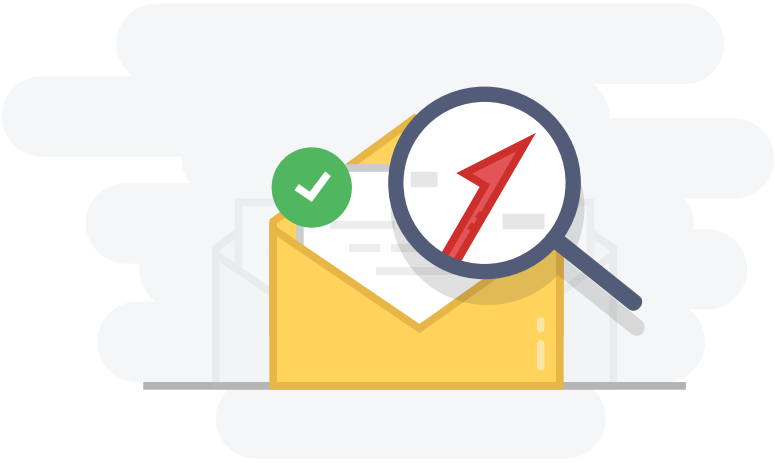

3. Ibid.

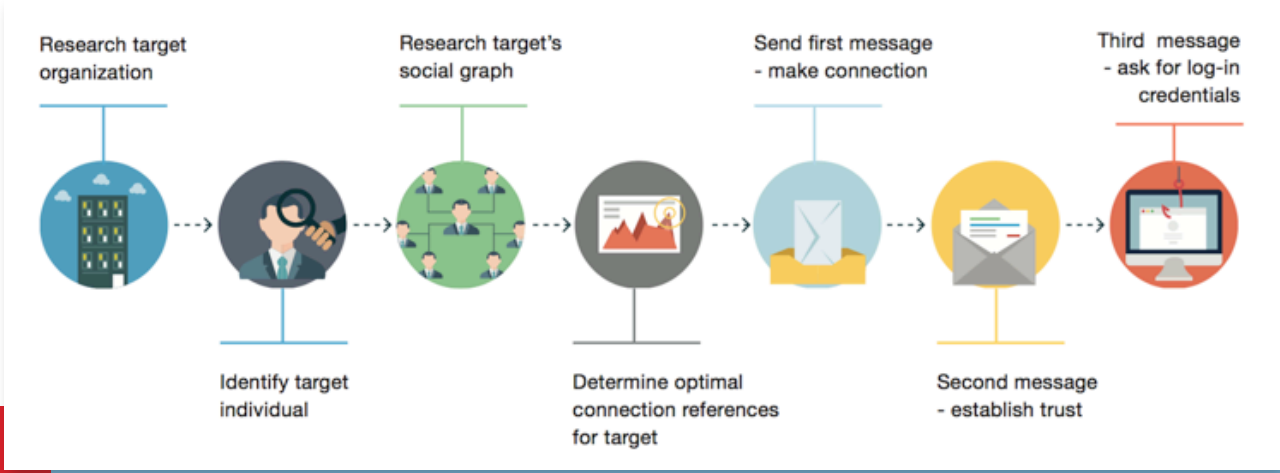

Figure 2 - The progression of a spear phishing attack, starting with research of the target organization and identification of a specific individual inside the organization, followed by a series of emails intended to build trust with the target.

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.

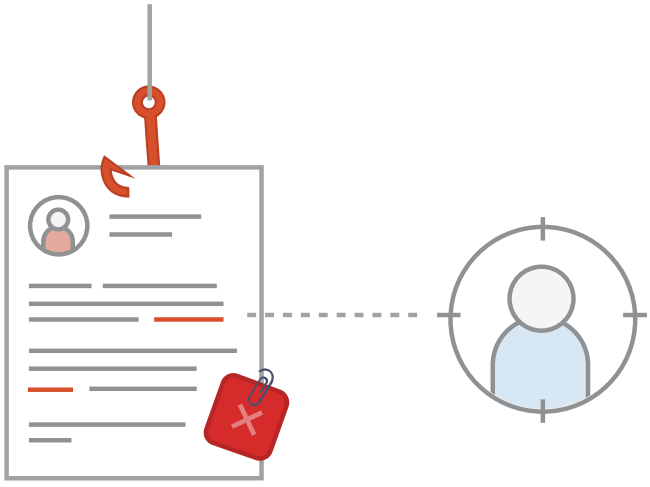

BUSINESS IMPACTS OF SPEAR PHISHING

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 enterprise $1.6 million.

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

Example: Sony Pictures

As a result of a 2014 hack, Sony Pictures experienced a high level of reputation damage, with private email exchanges between executives revealing embarrassing comments about famous people. The studio lost control of unreleased movies, which fell into the hands of digital pirates and had to be rushed into the marketplace. Millions in potential revenue was lost. Sony's brand suffered, which affected their valuation and ability to do business in Hollywood. Competitors gained insider knowledge of the goings-on at the studio. Finally, the company had direct costs such as $8 million to settle lawsuits with employees who were forced to protect their identities from theft after the incident.[4]

Spear-phishing attacks are often just the first part of a much larger hacking campaign. 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. And, of course, individual executive careers can suffer in the wake of such events.

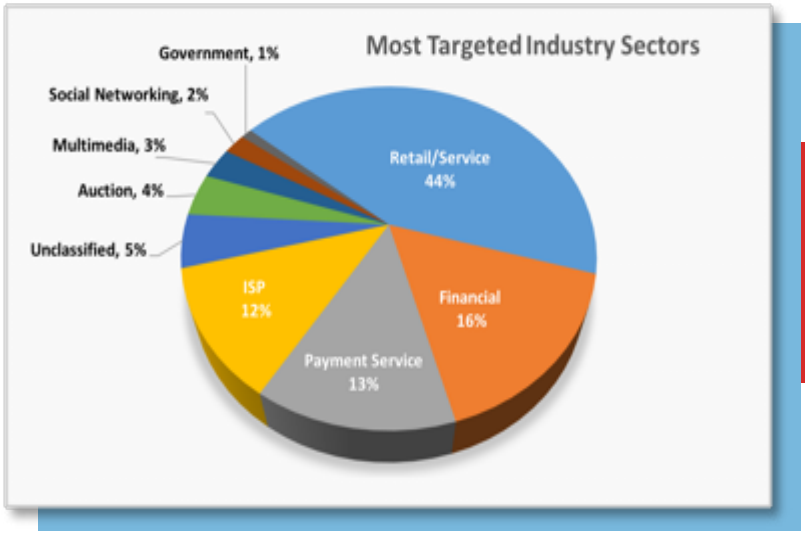

Figure 3 – Attacks by industry, 2nd quarter of 2016.

(Source: APWG Global Phishing Report 2Q 2016)

4. Brandom, Russell, "Sony Pictures will pay up to $8 million to settle hack lawsuit with employees." The Verge, October 20, 2015.

Risks by Industry

Financial Services: Financial firms must manage spear-phishing risks that can result in theft of inside trading information, personally identifiable information, credit card numbers, bank account information, and more. The impacts include financial loss, legal liability, and regulatory penalties.

Retail: As several large-scale hacks have shown, retailers are vulnerable to attacks that leak customer data — including credit cardholder 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, both of which have been built over many years and at great financial cost. Retailers also face an indirect risk from spear phishing, which is liability for fraudulent sales made with stolen credit card numbers. This may sound trivial, but it is not. Investigations are now revealing the existence of quite large-scale theft operations that steal merchandise from ecommerce sites and ship it abroad in bulk.[5]

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. These companies are part of an actual war — an undeclared cyber war that is raging despite its quiet, largely invisible profile. These companies especially tend to keep these types of breaches as quiet as possible, so there are likely many more successful attacks in this sector than are publicly known. The impact of sovereign cyber espionage may be hard to quantify, but 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. They face stiff financial and legal penalties for data breaches. There are reputation risks given the sensitive nature of leaked personal health information. As several large health insurers have discovered, there can also be significant costs associated with providing identity theft protection to tens of millions of policy holders who have had their names, addresses, and Social Security numbers compromised.

5. APWG Global Phishing Report 2Q 2016

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

The problem is that native Office 365 filtering systems, such as EOP and ATP, will not catch your typical spear-phishing email. 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 spear phishing. A one-off well-written email will generally get past most corporate spam filters, as they match no known signatures that signify "malevolent".

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

Spam-derived email security actually works *OK* for most mass-emailed phishing attempts, as the systems are able to block new variants of phishing attacks after the first few tens of thousands of emails are sent and the initial reports come filtering back. (Of course, this is of little comfort if one of your employees was the lucky recipient of one of the initial exploits ...)

... but completely useless at blocking sophisticated spear phishing emails.

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 enterprise needs a purpose-built email *security* system that will stop all types of email-borne threats—not just a glorified spam filter.

WHAT'S WRONG WITH JUST USING MICROSOFT EXCHANGE ONLINE PROTECTION (EOP)?

As stated, EOP can be moderately effective against *known* threats. The problem is that it is almost entirely helpless in fighting unknown threats … whether from zero-day code buried in an Excel file or a business email compromise (BEC) spear-phishing attack.

Here's what EOP is missing from a security perspective:[6]

- The ability to identify new and evolving threats for which it doesn't have a known signature.
- The ability to sandbox all attachments, such as .zip files.
- Real-time URL and page exploration to ensure links are safe and guard against time-bombed URLs.
- Robust spoofing detection.

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.

6. EOP is also missing some nice-to-haves such as effective graymail (low-priority email) classification, one-step unsubscribes, and archival capabilities.

Vendor Comparison

	Vade Secure	EOP	ATP
Known Threats			
Basic Signature-Based Protection for Known Threats			
Blacklisted Spam	✓	✓	✓
Blacklisted Malware	✓	✓	✓
Blacklisted Phishing	✓	✓	✓
UnKnown Threats			
Zero-Day Malware/Phishing/Spam Detection			
Basic IP/Domain/URL Reputation	✓	✓	✓
Basic Machine Learning & Heuristics	✓	✗	✓
Advanced Complex Reputational Analysis	✓	✗	✓
Anti Phishing Detection and Protection			
Real time URL Exploration	✓	✗	✗
Time-Bombed URL Protection	✓	✗	✓
Sensitive data request	✓	✗	✓
Exact Spoofing Detection	✓	✗	✓
Similar Sender Spoofing Detection	✓	✗	✗
End User Productivity			
Low Priority Email Classification	✓	! *	! *
One Step Unsubscribe	✓	✗	✗
Advanced Unsubscribe	✓	✗	✗

* limited to spam only, not other low-priority messages (newsletter, social notifications, etc.)

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. 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 500 million mailboxes worldwide, and processes billions of emails every day. We have a dominant position in Europe (over 90 percent of all email in France flows through our servers, for example) and a strong presence throughout North America and worldwide.

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, 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.

Artificial Intelligence plus Traditional Filters

Vade Secure's AI-based predictive email defense solution is supplemented by an extensive chain of additional protections, such as traditional signature-based spam filters, an extensive blacklist, and two complementary virus scanners.

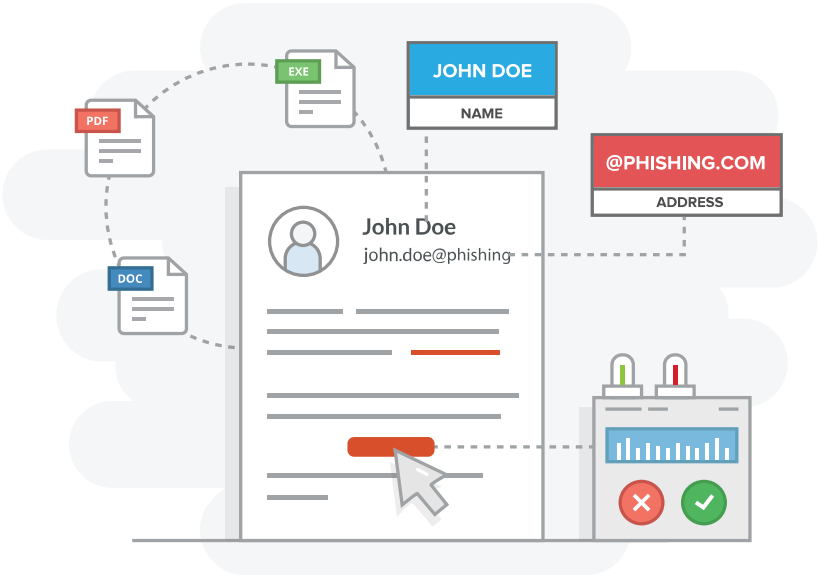

Vade Secure's Solution

The full Vade Secure email solution provides spam, graymail classification, and the most robust email security solution on the market.

- **Initial filtering:** Emails are analyzed for known phishing and malware signatures, including executable files. This quickly weeds out all spam and mass attacks.
- **Anti-Phishing:** Using smart patterns and machine learning algorithms, Vade Secure crawls the URL and webpage, following any redirections in order to reach the final page and determine whether it's fraudulent. Unlike most URL exploration engines, we explore the URL both when it first comes through the system and any time a user attempts to click on a link, thus defeating time-bombed URLs.
- **Anti-Spear Phishing:** To prevent one-off spear phishing or BEC attacks, 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. Upon detecting any anomalies, the solution displays a banner within the email alerting the user that the message might be malicious.
- **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 30 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.
- **Spam Control:** Vade Secure achieves a 99.99 percent catch rate with essentially a zero percent false-positive rate (<0.00001 percent).
- **Graymail Management:** Vade Secure automatically classifies low-priority message (e.g. newsletters, promotions, social notifications), while one-click safe unsubscribe easily eliminates unwanted communications, allowing users to have a cleaner inbox.

Deployment Options:

Deployment of Vade Secure is simple. Your admin can get the solution up and running with just a few clicks, simply by authorizing Vade Secure to access your Office 365 email flow and configuring your filtering policies. This allows you to take a layered approach that augments existing Office 365 security like EOP and ATP with Vade Secure's predictive email defense.

Spam, initial filtering, phishing, and malware detection are all immediately 100 percent effective upon deployment. There is no need for a learning period. Spear-phishing detection will launch at roughly 90 percent effectiveness, with maximum effectiveness being achieved after approximately two weeks as the system tunes itself to your organization's specific habits and styles.

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 company.

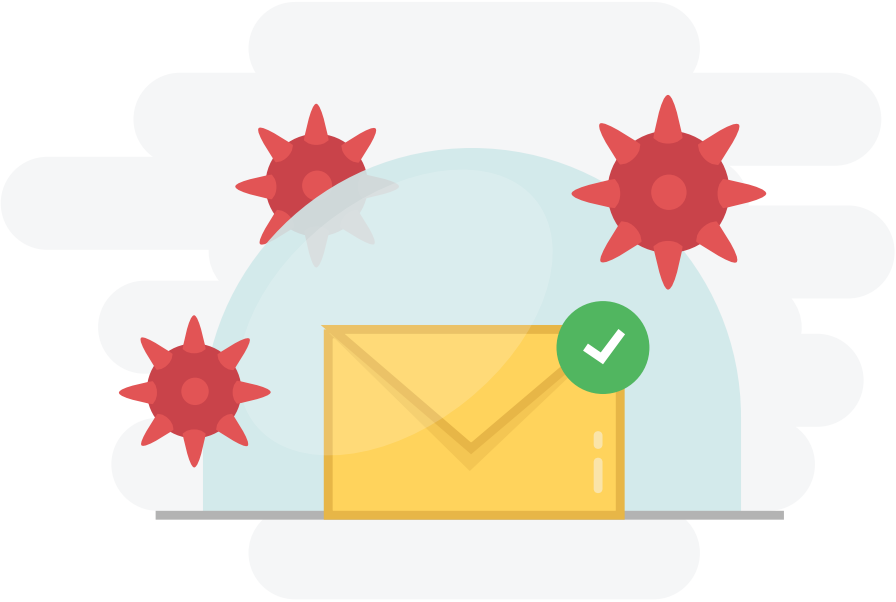

About Vade Secure

Vade Secure is the global leader in predictive email defense, protecting more than 500 million mailboxes in 76 countries. Our technology delivers best-in-class filtering accuracy through a layered approach that leverages artificial intelligence and machine learning in conjunction with smart patterns and human intelligence from our 24/7 follow-the-sun Security Operations Centers. We have a proven record of blocking unknown, small-wave and highly targeted phishing, spear phishing, and malware attacks from the very first email. In addition to email security, Vade Secure is a leader in more general email filtering, providing a comprehensive set of productivity-enhancing graymail management and safe unsubscribe tools. Our solutions are tailored for enterprises, as well as ISPs, telcos, OEMs, and hosting companies.

For more information about Vade Secure's AI-based anti-phishing solutions for Office 365, visit us at www.vadesecure.com

or call us at +1 (415) 745 3630.