SANS

Analyst Program

Whitepaper

2022 SANS Protects: Enterprise Email

Written by **Matt Bromiley**

July 2022

Abnormal

Introduction

Email is arguably one of the most important, if not *the* most important, communication mediums in modern business. Responsible for connecting billions of people and generating hundreds of billions in revenue, email is a quintessential part of communicating. Unfortunately, it is also one of the most prevalent and preferred attack vectors for adversaries. Verizon's 2022 Data Breach Investigation Report (DBIR) found that 82% of analyzed breaches involved the human element, with spearphishing being a significant contributing factor.[1]

Because organizations cannot disable email, they are constantly seeking technologies to help defend their email. From email validation and email origin tagging to advanced secure email gateways that can inspect and check every part of an email, security teams are continuously looking for ways to block malicious emails from entering their environment. However, as we see year after year, adversaries still find success by abusing email.

In this SANS Protects paper, we look at threats to enterprise email and ways that organizations can overcome or mitigate these dangers. Our SANS Protects papers focus on threats and mitigations, helping organizations consider elements of security that they should be implementing today. Our goal is to start or enhance the conversation about what assets are in your environment and the corresponding protections in place.

Given adversaries' success rates, organizations cannot simply ignore email threats. We recommend that organizations place email protections high on their list of security priorities, because adversaries find new techniques daily and security teams need to be able to keep up. Some of our key takeaways include:

- Utilizing modern capabilities, such as natural language processing (NLP) and machine learning (ML) aids, in detecting malicious "needles" in a haystack of millions.
- As a conversational medium, email language understanding and analytics can be used to identify statistically "out of place" emails that represent threats to users.
- Traditional defenses that look for malware and malicious links are still useful in capturing commodity and/or non-specialized phishing emails.
- User education will always be the final line of defense. Ensuring users know how to spot malicious emails can save a company millions of dollars or a high-profile intrusion.

We've seen plenty of organizations surprised when malicious emails get through, not knowing that oftentimes security features must be implemented or enabled. (They are not there by default!) It's easy to assume that every organization can implement the best email defenses or that those protections may already be in place. We recommend that, in lieu of advanced email defenses, organizations look at their email provider to see what sort of available defenses they can employ at zero cost. In many cases, preventive measures such as legacy protocol disabling or multifactor authentication (MFA) can be turned on, offering a layer of protection that can neutralize some of the most successful attacks.

[1] "2022 Data Breach Investigations Report (DBIR)," Verizon, www.verizon.com/business/resources/reports/2022/dbir/2022-data-breach-investigations-report-dbir.pdf [Registration required.]

Threats to Enterprise Email

Phishing (of All Kinds)

Phishing is more of a "parent" threat that encompasses multiple types of sub-threats. Of course, this taxonomy can vary depending on how an organization views phishing, response playbooks, and threat intelligence feeds, among other considerations. Table 1 identifies the classifications of phishing we use for the purposes of this whitepaper.

Within the phishing classifications described in Table 1 are subcategories of threats. For example, one type of phishing attack is business email compromise (BEC) or executive account compromise (EAC). This kind of attack may simultaneously target high-positioned individuals with specific instructions or target victims by pretending to be a high-value individual.

Table 1. Phishing Threats

Type	Description
Mass (bulk or commodity) phishing	Phishing attacks sent to multiple targets, which may include multiple recipients at the same or different organizations
Spearphishing	Targeted attacks sent to unique recipients, often with specific email content
Whaling	Targeted attacks against high-value or high net-worth victims, such as those in the C-suite
Smishing/vishing	Attacks that may break outside of email using text messages or video calls but have the same intent—to steal credentials or get victims to do something they would not normally do
Angler phishing	Social engineering attacks that may use social media posts to lure victims into providing credentials or downloading malware

BEC attacks are some of the most successful in cyberspace, yielding nearly $2.4 billion in adjusted losses in 2021, per the FBI's 2021 IC3 Report.[2] In the past decade, BEC attacks have resulted in more than $10 billion in losses, a number that grows every year. Even worse, adversaries continue to hone their craft with deep fake audio and usage of meeting platforms, clearly taking advantage of a now remote world.

Legacy Protocols/Authentication

Another top concern within email is the use of legacy/outdated protocols or authentication mechanisms. As mentioned previously, email has been around for decades. As such, the medium has developed in such a way that it has had to adapt to technological concerns over time. These adaptations include modern authentication mechanisms (such as MFA), secure protocols, and encryption. However, these advanced features are not always enabled by default.

Additionally, some software may require that legacy protocols or authentication mechanisms remain enabled to support things such as mobile email. This can create an issue for organizations that wish to modernize their email applications but struggle to convince all users to let go of their legacy applications. To help satisfy these users, IT administrators will leave legacy protocols open, hoping one day to eventually phase them out.

[2] "Internet Crime Report 2021," FBI, www.ic3.gov/Media/PDF/AnnualReport/2021_IC3Report.pdf

It is in this gap that some adversaries find success. We have seen attacks take place where adversaries will intentionally, via stolen credentials, force a single-factor authentication to an email account. Even users who have MFA enabled can have their authentication subverted if a single-factor protocol is enabled. This allows attackers to steal a username and password to access an email account, despite MFA being in place.

Malware Delivery

Perhaps one of the most common uses of email from an adversarial perspective is delivery of malware to a victim user(s). This is one of the most tried-and-true delivery mechanisms, especially if you can evade email defenses. Users trust that an email that shows up in their inbox has already been vetted and scanned, and thus they trust the contents implicitly (without any additional user education, of course). Some adversaries may have a conversation with the user to gain that trust before sending the malware.

This approach proves to be such a successful delivery mechanism for adversaries because users tend to send everything they can via email. Whether it's a Word document, PDF, Excel spreadsheet, or even sometimes raw code, it is not out of the ordinary for someone to send a file from one user to another. With trust in place, users are accustomed to the routine act of downloading and opening files. This can lead to users opening files, allowing code to run, and providing attackers with the access they need to gain further entry into an environment.

Email Platform Compromise

Another critical issue when it comes to email security is the security of the platform itself. Remember, email platforms are pieces of software running on servers. As such, they are susceptible to vulnerabilities, code flaws, OS exploits, and a slew of other attacks that adversaries may launch against systems in an organization's environment (including things such as credential harvesting, privilege escalation, and/or lateral movement).

Email platforms often store email data, user data, and email configuration information on the servers themselves. Because of the valuable information these attributes contain, they place email servers in the crosshairs of many adversaries. Furthermore, email platforms are designed to be internet-facing. After all, email must be sent and received across the internet! A vulnerable email server can create yet other entry vector for adversaries, without anyone ever needing to send an email.

We saw this come to light in March/April 2021, when a slew of vulnerabilities was announced for Microsoft Exchange Server on premises. Via specially crafted packets, adversaries could exploit an internet-facing email server to gain initial access into an environment and/or download user details such as usernames and passwords.

Outbound Threats

Remembering that email is a bidirectional medium, threats "to" email also can come from within. These aptly named *outbound threats* encapsulate user mistakes and malicious insider intent, both of which can lead to sensitive data leakage and/or oversharing or exfiltration of data. Table 2 lists examples of key outbound threats.

Spam

Last, but certainly not least, is the issue of spam email. This may not necessarily be the largest threat to organizations; however, it does represent a nuisance for both users and email detection platforms. So much spam is sent daily that it can overload users who are trying to discern between what's good and what's bad. It may also overload email platforms doing their best to separate good from bad in real time for hundreds or thousands of users simultaneously.

Table 2. Outbound Threats

Threat	Risk
Misaddressed emails or incorrect attachments	Users can inadvertently send sensitive or damaging information outside of the organization and into an untrusted space.
Accidental distribution list	Data intended for one or two recipients can accidentally be sent to an entire distribution list, putting data at risk of oversharing and/or leakage.
Unencrypted sensitive data	Emails and email attachments that are unencrypted may be intercepted and read by malicious parties who have access to an account or the email transfer process.
Data exfiltration	Email attachments allow for data to be sent "out" of an organization. This can be accidental by an unknowing employee or done intentionally by a malicious insider. Either way, traffic from inside-out may not be subject to the same filters and thus represents a potentially unmonitored exfiltration route.

As with commodity malware, spam can be a nuisance threat that eventually leads to a significant compromise. For example, consider a spam email that attempts to steal a user's credentials. Sent to thousands or millions of people, the attacker or spam email author may be unable to utilize every single credential they have acquired. However, this database full of stolen credentials, collected over time, may end up being sold multiple times to multiple adversaries and subsequently used to enter an environment via an open port or remote access tool.

Protecting Enterprise Email

Protecting email is not as simple as *stopping* email. The problem is that the same medium that organizations use to conduct transactions is also used by adversaries to deliver malware. Because of this, we must find ways to correctly parse and interpret emails to understand what is legitimate and what is malicious. Of course, this is often much easier said than done.

Furthermore, when we consider solutions to protect email, we also must consider real-time actions. We cannot enter a state where emails are delayed for the sake of security, because unfortunately that might have adverse effects on the business. Thus, we look for features within solutions that can take place in real time, or near real time, and still offer a high efficacy rate in blocking malicious emails.

Natural Language Processing

One of the first features to look for in an email platform is natural language processing (NLP) or natural language understanding (NLU). Email is a conversational medium. This often means that emails contain large blocks of text, punctuation, attachments, and other items that, to a reader, would appear normal. It is not uncommon for an organization to email customers, send attachments, request financial details, or discuss employees or future business plans via email. Unfortunately, this creates a common dictionary of words shared among businesses and attackers—confusing readers who aren't on the lookout. For example, the phrases "please send money" or "attached is the document you requested" are not inherently malicious based on those words alone.

Understanding the natural language of an email is a complex but useful feature that can be used to analyze the email and its context and metadata. It also can provide the user or security teams with a decision as to whether the conversation appears to be a legitimate transaction or not.

The power of NLU is bolstered when additional data points are available for the decision process (such as the features we discuss later in this paper). In fact, NLU is most effective when other technological filters have done their job. For example, an email that passes authentication checks, does not contain any malware or malicious links, and does not trigger any heuristic analysis faults would be a great candidate for NLP to determine whether that email is malicious or not.

Email Heuristics and Graph Analysis

Another valuable feature to look for in an email security platform is email heuristics, or analysis of email metadata. When we think of email heuristics, we should consider questions such as:

- How many times have the parties conversed?
- What is the context of a normal conversation between these parties?
- Are the parties internal or external?
- Do their emails normally involve attachments? If so, what are the attachments?

Simple email heuristic analysis can help us distinguish between malicious emails and good emails, even if attackers take the time to spoof superficial metadata. This is often how users fall victim to phishing emails. Attackers can spoof certain fields of an email, such as the *To:* or *From:* fields, which subsequently makes the victim believe that the sender is legitimate. (We will address some technologies to mitigate this in the next section.)

While it can be useful to stop a certain level of malicious emails, email heuristics and graph analysis can easily fail to capture malicious emails from a third party (such as suppliers) that may have been compromised and is being used to launch an attack on trusted relationships. This is why good email defenses rely on multiple technologies working together, avoiding a single point of failure.

Email Detection/Prevention

Combining a multitude of technologies, an email detection or prevention system, sometimes also known as *secure email gateway (SEG)*, is an all-in-one platform that can be used to analyze, detect, and prevent malicious emails from entering an environment. These platforms often find success by combining a series of technological concepts, many of which we've discussed in this paper. The goal of such a platform is to create a seamless experience for users while giving security teams absolute control over emails coming in and out of the environment.

A secure email gateway (SEG) can protect against known threats and is an excellent solution to help protect enterprise email. However, sophisticated attacks can still bypass a SEG. Consider augmenting it with an integrated cloud email security (ICES) solution, which can hook directly into cloud email solutions and defend against advanced threats.

SPF/DKIM/DMARC

Perhaps some of the oldest technology recommendations in this whitepaper are to implement some basic email protection mechanisms that have been around for years. Used to enable email authentication and genuineness, these are simple technological capabilities that are often built into most email platforms and can be enabled—if they have not already been enabled.

- **Sender Policy Framework (SPF)—**A protocol that confirms an IP address is allowed to send an email on behalf of a particular domain. It can be used to prevent spoofing.
- **DomainKeys Identified Mail (DKIM)—**An ID that can validate a sender's identity and determine whether the email was changed during transit. This also helps confirm an email is authentic and not spoofed.
- **Domain-based Message Authentication, Reporting, and Conformance (DMARC)—** Built on top of SPF and DKIM, DMARC is a protocol for email authentication, policy, and reporting. DMARC addresses malicious emails by matching up the *From:* field with the sending domain and provides the domain owner the ability to determine how to handle unauthenticated emails.

While these are commonplace technologies, they are not always enabled in email platforms—not even in cloud email platforms. However, while they are not always trivial to implement, we highly recommend that security teams, IT administrators, and email server administrators work to enable these authentication mechanisms. They can be used to protect both your users and the recipients of emails they send.

User Education

Finally, and perhaps one of the most important lines of defense in protecting against email threats, is user education. This is not necessarily a technology that can be implemented (unlike previous recommendations), but more of a combination strategy that security teams should employ inside of their organization with technology and training. Adversaries are counting on your users not noticing that their requests or their emails are malicious. They are counting on users clicking malicious links or showing interest in the attached items.

One of the best investments organizations can make in their security defenses is enabling and empowering their users to detect and report phishing emails. This can be achieved via a variety of methods, including:

- Conducting phishing tests or simulations, which determine if users click on malicious emails and just how far down the path they go
- Including one-click email defenses in user email platforms or applications, to simplify the reporting process
- Rewarding users and/or the organization for effectively stopping a phishing attack

The numbers prove that security awareness training works. A report from Osterman Research in August 2019 found that smaller organizations (as many as 1,000 employees) can achieve an ROI of 69% from security awareness training, while larger organizations (more than 1,000 employees) can achieve an ROI of 52%.[3]

Closing Thoughts

As a communication medium, email has been around for more than 50 years. The first email, sent in 1971, contained simple text and was sent from a user to himself. Since then, the various technologies surrounding email to secure and protect it have changed, but the basic concept of sending a message between at least two parties remains the same. And unfortunately, within this medium, attackers have found an easy way to enter an environment. As security teams find ways to mitigate their environment against email threats, they often turn to myriad technologies to help defend one of their front lines.

In this SANS Protects paper, we looked at some of the common threats to enterprise email and defenses that can be implemented to mitigate these threats. In many cases, threats depend on adversaries abusing the conversational nature of email. However, we simply cannot just turn email off. Instead, we examined technology implementations and user education techniques that can be used to mitigate email threats. With the right combination, organizations of any size can significantly lower their chances of falling victim to an email-based attack.

[3] "The ROI of Security Awareness Training," Osterman Research, www.infosecinstitute.com/wp-content/uploads/2021/03/IQ-Whitepaper-The-ROI-of-Security-Awareness-Training.pdf

However, even the best-laid defenses do not mean security teams can simply assume attacks will no longer happen. Adversaries change their tactics, techniques, and procedures (TTPs) all the time, finding new ways to send malware into an organization that may have robust email defenses in place. Time and time again, we see cases where adversaries befriend users prior to launching their attack, to gain trust and ensure that the victim will complete the actions requested of them. Therefore, it is so important to remember that users are your last line of defense: Training your users is just as important as training your security technology—they might save the company.

Sponsor

SANS would like to thank this paper's sponsor:

Abnormal