Office 365 helps secure Microsoft from modern phishing campaigns

Contents

Office 365 helps secure Microsoft from modern phishing campaigns

Cybersecurity is a critical issue at Microsoft, as it is for organizations everywhere. Microsoft processes more than 400 billion emails each month and blocks 10 million spam and malicious email messages *every minute* to help protect our customers from malicious emails.

Phishing attacks are designed to trick people into sharing credentials or personal financial information. According to an FBI report, phishing attacks are increasing, costing an estimated $5 billion in compromised business email since 2013. In 2016, the Anti-phishing Working Group (APWG), which Microsoft is a member of, saw more than 255,000 unique phishing campaigns with attacks on over 600 brands. In a six-month period in 2017, there were over 800 million phishing mails flagged in Office 365. Protecting against phishing is a persistent need for most, if not all, enterprises.

IT organizations that support everything from small businesses to global enterprises, including Microsoft Core Services Engineering and Operations (CSEO, formerly called Microsoft IT), rely on Office 365 mail services. Microsoft has heavily invested in sophisticated anti-phishing technologies for many years to help protect our customers and our employees from constantly evolving, increasingly sophisticated, and often targeted phishing campaigns.

Office 365 Exchange Online Protection (EOP) and Office 365 Advanced Threat Protection (ATP) work in near real-time to protect against phishing threats and safeguard data and intellectual property. EOP provides advanced security and reliability to help protect information and eliminate known threats before they reach the corporate firewall. Office 365 ATP further protects mailboxes against new, sophisticated attacks by expanding protections against unsafe attachments and malicious links. It complements the security features of EOP to provide better protection against zero-day, advanced, and targeted phishing campaigns.

Evolving protection in a changing threat landscape

The Microsoft approach to protection against phishing in Office 365 is dynamic and robust, and evolves with the strategies and tactics used by attackers. Shared signals across Office 365, Windows, Azure, the Microsoft Intelligent Security Graph and first- and third-party antivirus (AV) engines make Microsoft uniquely positioned to protect against phishing attacks.

Every month, Microsoft combines intelligence from 400 billion emails analyzed by Office 365, over 1 billion Windows devices, and 450 billion user authentications from Azure Active Directory (Azure AD), as well as signals from hundreds of other services and properties. This breadth and depth of security signals and data is used to power rich machine learning, AI algorithms, and heuristic algorithms that fuel the creation of new detection technologies.

Rich data sources combined with cloud intelligence in Office 365 is helping surface more actionable insight that helps our security administrators manage security and compliance within Microsoft.

Phishing is a complex threat

The phishing landscape has many types of attacks, ranging from basic scams like emails requesting information or financial transactions from foreign dignitaries, to highly sophisticated and targeted spear-phishing campaigns that impersonate a brand or a well-known person. Industry research suggests that 91 percent of security breaches originate from phishing or spear-phishing.

In the past, phishing emails weren't carefully constructed or targeted. Today, however, phishing cyberattacks come from a criminal industry that includes companies, crime organizations, and even nation-states. These threat actors employ research and surveillance teams to:

- **Steal personal information.** An attacker goes after personal account details, passwords, credit card numbers, or other sensitive information.
- **Steal credentials.** An attacker obtains a person's username and password—for example, credentials that an employee uses to sign in at work.
- **Identity theft.** An attacker impersonates a person or company with compromised credentials and other information.
- **Harvest confidential information**. Once an attacker gains access to a device or network, they can quietly collect proprietary information or inspect the environment for vulnerabilities. They can look for information that will give them even more targeted spear phishing methods.

As shown below, the phishing attack spectrum can range from broad to targeted, using a complex variety of lures.

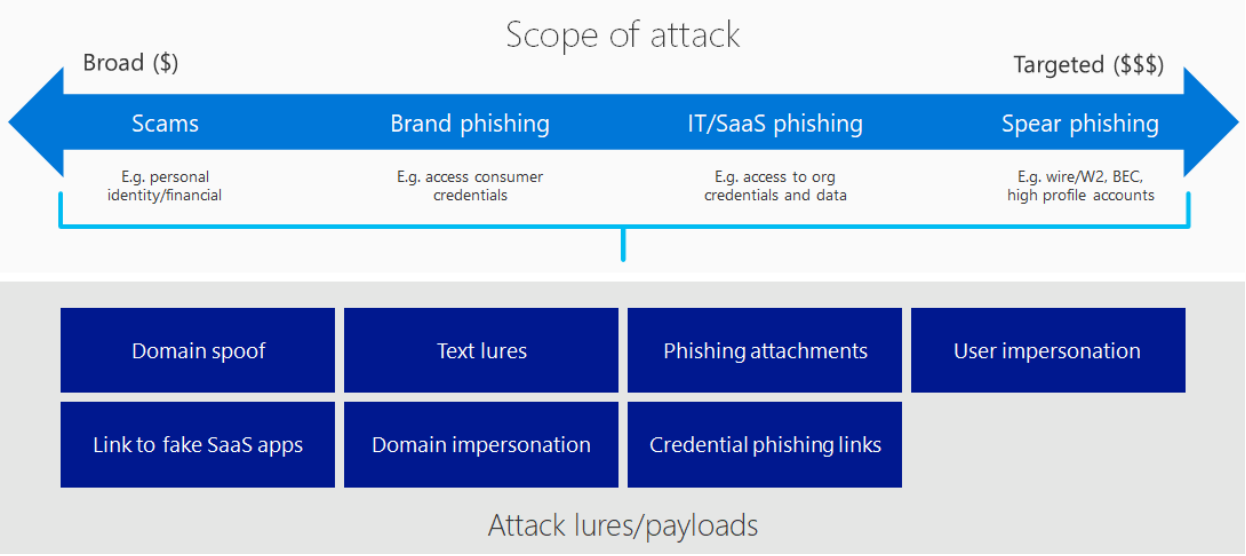

Figure 1. Phishing attack spectrum and lures

Broad phishing casts a wide net. In most cases, these attacks are basic scams that target people and seek personal information to compromise the user's identity or financial information.

Brand phishing is designed to trick people into signing into a spoofed site or clicking on a link to enter, confirm, or reset their password. These emails typically go after consumer account credentials for things like peer-to-peer payment systems, social media accounts, or even e-commerce accounts. The goal is to gain access to whatever is available in the compromised service and to harvest credentials. Since people often use the same credentials across multiple services, once a bad actor gets user credentials, they can access that user's accounts across several platforms.

IT and software as a service (SaaS) phishing is more targeted and more financially damaging. These campaigns target specific organizations, typically using a message that spoofs their IT organization or a popular SaaS app that a company runs their workflows on. These attacks are designed to gain access to the organization's credentials and then to use them to laterally compromise the organization and gain further access to corporate data.

Spear-phishing poses a special challenge because it is very targeted toward specific individuals or roles within an organization and can perpetrate the most financially damaging attacks, like W-2 fraud, wire fraud, compromising a high-value, high-profile account. Because of the proprietary information that executives have access to, they are often targets in spear-phishing attacks known as *whaling*.

Layered protection against a variety of tactics

Phishing campaigns can use combinations of lures to deceive recipients. To help address the variety of threats, Office 365 EOP, Office 365 ATP, Cloud Application Security, and Office 365 Threat Intelligence work together to offer layered protection with time of delivery, time of click, and post-delivery protection.

Figure 2 illustrates our different anti-phishing technologies within the context of mail flow.

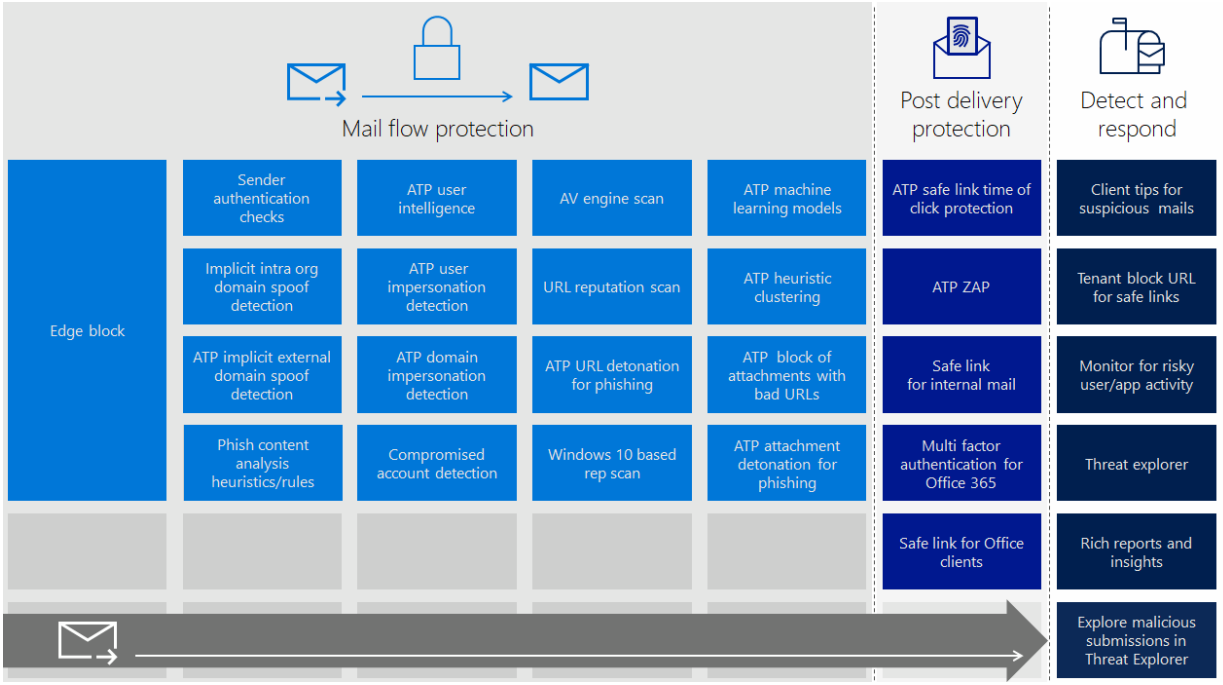

Figure 2. Multi-layered phishing protections in Office 365 and Exchange Online Protection ATP

Office 365 applies various anti-phishing technologies before email arrives in an inbox. EOP includes an edge block (IP and sender reputation), anti-spoof capabilities, authentication checks for both external and internal emails, link reputation lists, and sophisticated multi-engine reputation and AV filters. Office 365 ATP offers advanced algorithms that can detect user impersonation, domain impersonation, and implicit domain spoofing. It also features URL detonation, attachment detonation and blocking, user intelligence, reputation scans, heuristic clustering, and machine learning models that constantly improve phishing identification. This holistic, in-depth security layer prevents most phishing emails from ever arriving in a mailbox.

Because phishing threats are constantly evolving, some highly sophisticated and advanced phishing mail can make it to an inbox. For these more advanced phishing attacks, CSEO has enabled multifactor authentication capabilities offered for Office 365. Additionally, Office 365 ATP offers Safe links time-of-click protection. Office 365 security also offers Zero-hour Auto Purge (ZAP). ZAP continuously monitors email and will move a malicious message to the junk folder *even after* it has been delivered. With ZAP, Office 365 can help ensure that if a malicious email makes it through the pre-delivery scan and is later identified as phishing, it will be removed.

Use cases

This section illustrates how the layers of protection work when applied to some of the common phishing scenarios that we see in our environment.

Phishing emails that spoof a domain

Spoofing is a common way for threat actors to send phishing mails. When a sender spoofs an email address, the message appears to be sent from a domain but originates from an unauthorized source. If the recipient assumes that the mail came from a real domain, they might end up clicking on a malicious link or divulging private information to the attacker.

Explicit authentication checks

Sender authentication protection helps flag spoofed messages before they reach the user's inbox. An email's digital envelope contains information that Office 365 EOP authentication scans to determine if a sender is who they claim to be. The scans look at the Sender Policy Framework (SPF) to authenticate mail senders, the DomainKeys Identified Mail (DKIM) signature to determine if it originates from the domain, and the Domain Message Authentication Reporting & Conformance (DMARC) record associated with the sender's domain.

DKIM adds a digital signature to all outbound mail from within an organization. That digital signature can help confirm that the mail is actually coming from the organization. SPF record validates the origin of email messages by verifying the IP address of the sender against the owner of the sending domain. SPF also checks if a sender is permitted to send on behalf of a domain. If the sender is not permitted to do so, and the email fails the SPF check, DMARC helps us decide what to do with the message—whether it should be quarantined in the junk or spam folder, or rejected.

Advanced implicit authentication checks

In addition to explicit authentication checks, Office 365 also uses an additional anti-spoofing layer which implicitly infers authentication for mail traffic for domains that have not fully configured SPF, DKIM, DMARC, by learning from historical traffic patterns from those domains. Office 365 does this both for the receiving organization when someone tries to spoof their own domain, and is in the process of rolling out similar checks for all external-sending domains.

Allowing legitimate spoof email in O365 Security & Compliance Center

Office 365 has built-in anti-spoofing protection designed to detect legitimate spoofing—when someone needs to send email on behalf of someone else—while shielding the organization from illegitimate ones. Legitimate spoofing, for example could be when you have hired an external company to create and send out advertising or product updates on your behalf, or when an assistant regularly needs to send email on behalf of an executive.

Normally, Office 365 treats these spoofed messages as spam. You can prevent legitimate spoofed mail from being blocked by setting up spoof filters in the Security & Compliance Center.

Spoof intelligence

Sometimes, Office 365 does not have enough historical information to determine whether a spoof is legitimate or malicious. This can happen when a new sender starts sending email as someone else without the proper SPF, DKIM, or DMARC configuration, or if the volume of email is too small to generate a positive reputation. Office 365 ATP includes spoof intelligence, which can be accessed through the Anti-spam settings page in the Office 365 Security & Compliance Center. With Spoof Intelligence, our analysts can review all senders who are spoofing our organization and then choose to allow or block the sender and better manage false-positive cases.

For more information, see Learn more about spoof intelligence.

Phishing emails that impersonate a user or a domain

As phishing awareness increases, employees have become better at recognizing some of the more common phishing scams, such as those claiming to be from a lottery they didn't enter—or a foreign official trying to move money out of their country. Some phishing attempts are more difficult to discern because they use visual tricks to make email look like it is from someone you know or from a partner or service provider you work with.

The attacker relies on visual tricks in the display name or the domain name of the sender's email address to make it look like someone you know or like a familiar organization's domain. A common example is a Business email compromise, where the attacker uses these tricks to make the email address look similar to the CEO of the organization. The mail might go to the CFO or another high-ranking officer, and will ask the person to take an urgent action.

Office 365 ATP anti-impersonation settings

The new Office 365 ATP anti-phishing policy allows us to configure both user impersonation and domain impersonation detection settings. Our administrators can specify the users and key domains that are likely to get impersonated and manage the policy action like junk the mail or quarantine it.

To learn more about configuring impersonation detection in the new anti-phishing policy, see Set up Office 365 ATP anti-phishing policies.

Office 365 Mailbox intelligence

Behind the scenes, Office 365 builds user-level mailbox intelligence that figures out the strength of relationships between senders and receivers. Mailbox intelligence detects when an email is the first message received from a sender, and uses that information to determine the likelihood of it being phishing; then it runs required anti-impersonation checks. Two people that send a lot of mail back and forth have a stronger relationship, whereas email from a first-time sender indicates a weaker relationship. Since many phishing emails come from first-time senders, CSEO can refine policies for mail delivery based on our level of confidence and the strength of relationships.

Safety tips for impersonation

Safety tips can be used to prompt people when Office 365 detects suspicious or phishing email. For example, in the impersonation policy, CSEO can turn on the safety tip that warns users when they receive an email from an impersonator.

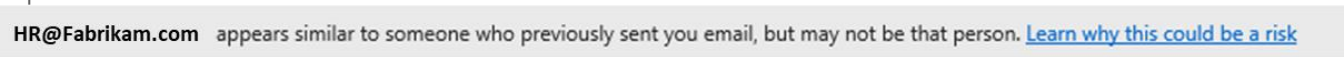

Figure 3: Client mail tip

Social engineering and emailed links to phishing sites

Some phishing mails ask a user to click on a URL to change their password, sign in to see an offer, or to access something in a secure message center. While the visible link may look legitimate, the embedded link in the mail can go to a duplicated site or to a sign-in page that intends to capture their user name and password. As illustrated below, users will soon be able to hover over an embedded link in the body of an email to inspect its URL.

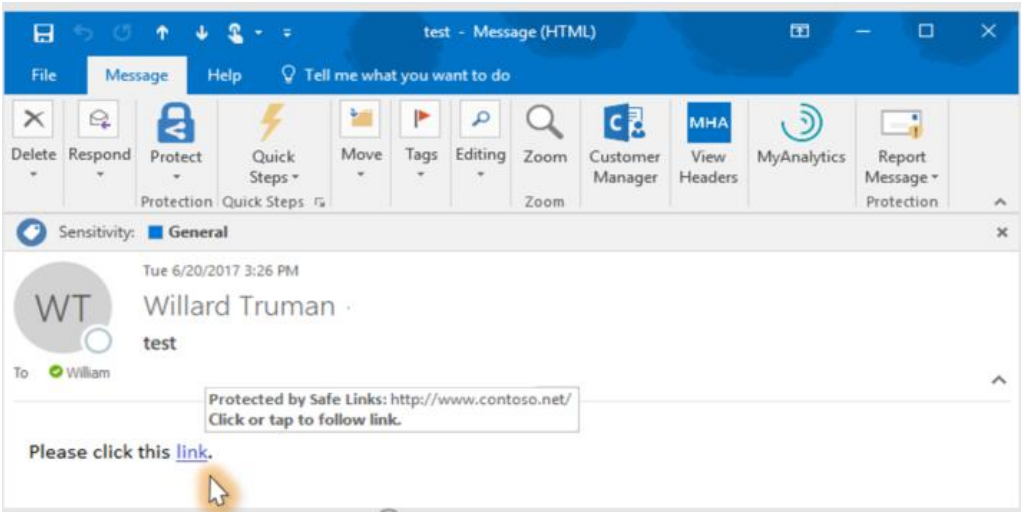

Figure 4. Hovering over a Safe Links-protected URL

URL scanning, detonation, and Safe Links

For our environment, we have created Safe Links policies that check and block malicious URLs in email that Microsoft employees get from outside the company. URLs listed as malicious in Office 365 reputation scans will be marked as spam and will be blocked when the user clicks on them. If Office 365 doesn't block the mail based on any of the other

scans, Safe Links will open and analyze the link and site content, within a virtual detonation environment, to check for the presence of a lure before continuing to the website. Safe Links is updated with the knowledge gained through detonation. Office 365 allows us to configure policies to block malicious links entirely, or we can notify users that we don't know or don't trust the link, and they can choose to proceed if they have confidence in the link.

To learn more about creating Safe Links policies, see Set up Office 365 ATP Safe Links policies.

Machine learning

Machine learning and AI constantly improve the way Office 365 detects phishing emails. Office 365 machine learning models that look at various email properties such as the email header, the email body, and any links to detect phishing. Office 365 can follow links to a landing page and use machine learning to see if the landing page has any potential phishing lures.

Attachment-based phishing email

When phishing messages include an attachment, Office 365 will either block the message or move the attachment to a virtual detonation environment. Information about the mail and the attachment are used to inform reputation scanning signals and our machine learning models.

Safe attachments file detonation

At CSEO, we use Safe Attachments to configure policies that block phishing lures and malware in attachments. All attachments without a known signature are routed to a virtual detonation environment, where Office 365 ATP applies behavioral analysis and machine learning techniques to detect malicious activity. If no malicious activity is detected, the attachment is released for delivery. Safe Attachments support the ability to check files and files behind links. URLs within files can also be analyzed.

After detonation, Office 365 ATP updates its file reputation store so that any subsequent emails containing a previously detonated and flagged file is blocked by our EOP service. ATP Safe Attachments policies can be applied to specific people, groups, or your entire domain. To learn more, see Set up ATP safe attachments policies in Office 365.

Mail sent from compromised accounts

It can be difficult to detect a phishing or malicious email from a compromised account. It will not have signs of spoofing or impersonation, and Office 365 might not immediately flag it unless it includes a link or an attachment that has a phishing or malware signature. Office 365 and machine learning have helped us create scans to catch behaviors that indicate a compromised account, and our analysts can quickly investigate and respond using Threat Intelligence.

Impossible time travel

Office 365 uses Azure AD for account logins. If an account is logged into from another geographic region without enough time for the account holder to travel, or if there are log-ins from two locations at once, machine learning models will detect it and our security analysts will investigate to find out if the user has been compromised.

Cloud Application Security

Comprised credentials can do more than send mail—at CSEO, we have implemented Cloud Application Security (CAS) to help manage and limit cloud app access based on conditions and session context, including user identity, device, and location.

Protecting the mailbox after delivery

Zero-hour Auto Purge

ZAP continuously monitors for new spam, malicious attachments, or phishing URLs, and will move an email to the junk folder if it is malicious—even if it initially made it through the email protection stack. Employees are protected from malicious emails continuously.

Time-of-click protection

For emails that have been delivered to an inbox, Office 365 ATP time-of-click protection with Safe Links will check the link's reputation again before it allows the browser to open the page.

Detection and response

At CSEO, we assume that a small percentage of phishing attacks may get through. With Office 365, we can quickly respond to breaches, mitigate their impacts, and play a role in helping improve our detection strategies to prevent future attacks.

Office 365 Threat Intelligence

Office 365 Threat Intelligence is a new dashboard in the Office 365 Security & Compliance Center, shown in Figure 4. It uses the Microsoft Intelligent Security Graph to analyze billions of data points from global datacenters, Office clients, email, user authentications, and other incidents that affect the Office 365 ecosystem—as well as signals from our Windows and Azure ecosystems to get insight about attacks.

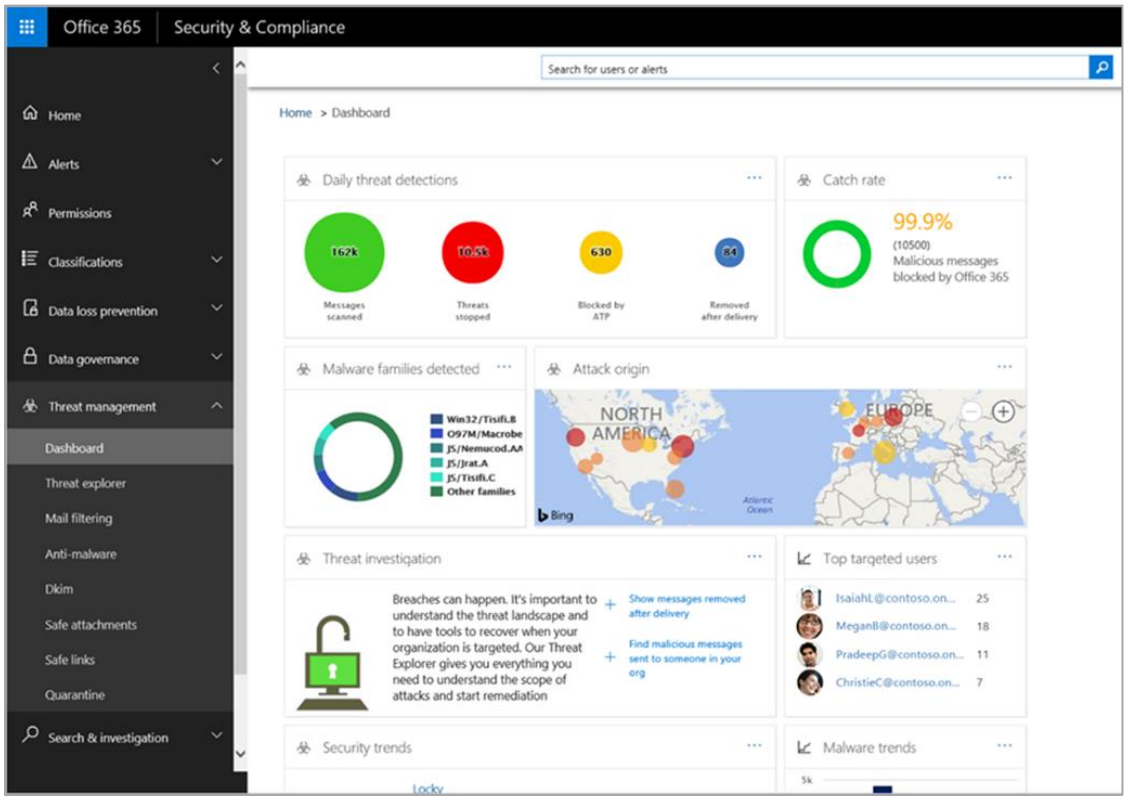

Figure 5. The Office 365 Threat Intelligence Dashboard provides visibility into the global threat landscape

Threat Intelligence works with other Office 365 security features, like EOP and ATP, so at CSEO we can see a wealth of information, including the top-targeted people and roles, the frequency and scope of an attack, and any available

security recommendations. We can also use features like Content Search in Office 365 Security & Compliance Center to see the body of malicious email and get full context for further analysis.

Threat Explorer

Using Threat Explorer, included in Threat Intelligence, we can view and analyze information about malware inside and outside our environment, including breach information. From Threat Explorer we can:

- Triage and investigate user-submitted emails that bypassed EOP and ATP.
- Gather details on active phishing attacks such as sender, recipient, source IP address, file hashes, subject lines, or URL links to identify affected users and impact on our environment.
- Prevent users from interacting with malicious emails that made it to their inbox by taking specific actions such as moving them to junk, deleting the email, or deleting attachments.
- Search for indicators of current and emerging email threats across our environment to determine impact and identify areas that require response.

Reported phishing

At CSEO, we receive reports of phishing from the helpdesk and through the **Report Message** add-in for Outlook. Working backward from the reports we receive, we use Threat Explorer and other security and compliance views to analyze the message, investigate the scope of the attack, and collect details about senders, attachments, and links. They are triaged, prioritized, and escalated for proper mitigation.

We also receive alerts in our security information and event management (SIEM) system. Working back from those alerts, we identify the entry point and, in most cases, it is the result of a phishing attack. The SIEM, in conjunction with Threat Explorer, has proven to be a powerful combination that enables CSEO to detect and respond to active attacks faster, from detection through response and remediation.

Improving detection and response capabilities

From a security and incident response perspective, CSEO security and monitoring processes are structured like most other enterprise-level security operations centers. We are in a unique position to help influence the direction of Microsoft products, based on our experience. Detailed information from phishing attacks that we investigate, or that are reported by employees, are presented back to the Office 365 product group to improve and evolve security capabilities.

Things to consider as you begin configuring Office 365 for phishing detection and response

If your organization has just begun to use Office 365 EOP and Office 365 ATP to protect from phishing, here are some things to consider, based on our experience at CSEO.

- Enable Safe Links and Safe Attachments to help protect your people. Both can be implemented with minimal disruption to email delivery.
- Adjust confidence thresholds for anti-phishing machine learning models. Office 365 allows you to tune the aggressiveness of the verdicts that machine learning models deliver. You can adjust confidence thresholds for specific users or user groups. For example, you can set policies to ensure that your executive's accounts have a low tolerance for phishing. So even messages that Office 365 marks with low to medium confidence, the phishing threshold can be adjusted so that the email is marked as high-confidence phishing. Adjusting the confidence threshold is an effective way to provide an extra level of protection.
- Enable multifactor authentication for your users. Your employees, like many of ours, probably use a combination of the same password and email address which can be risky, especially when they access resources outside of your organization. To help secure your employee's identities when they access mail from outside the corporation, consider enabling multifactor authentication for all your Office 365 users.

Note: For more information about using multifactor authentication in Office 365, see Set up multifactor authentication for Office 365 users. After you enable multifactor authentication on your tenant, your users can refer to Set up 2-step verification for Office 365 to set up their second sign-in method.

- Configure and enforce Domain Authentication. To round out your protection, set up SPF, set up DKIM, and set up DMARC to help protect your own domain from being spoofed.
- Disable SMTP-based login into Office 365 accounts. In today's modern workplace, most organizations don't need their users to connect and send email using SMTP protocol (such as old imap or pop3 clients) and it should not be broadly enabled. Office 365 now supports disabling SMTP-based login for an organization. Consider disabling this to further block a compromised user's ability to send internal phishing mails
- Increase insight with phishing reports and URL Threat Explorer. Gain rich reporting and URL tracking capabilities as well as insight into who is being targeted in your organization and the category of attacks you are facing. ATP reporting allows you to investigate messages that have been blocked because of an unknown virus or malware while URL trace capability allows you to track individual malicious links that have been clicked.

You will have better visibility into who is getting phished. Knowing what people or groups are receiving more phishing attacks and those who appear more vulnerable to risks, based on their computer use behavior, will help you refine policies and thresholds that can help reduce risk.

- Install the new junk mail reporting add-in, Report Message, to report phishing emails that are missed. In addition to improving Office 365 phishing filters, the reports can be used by your security and monitoring team in the Security and Compliance console. With those reports, you can quickly investigate the scope of an attack and take action. If the same mail has been delivered to other mailboxes, you can clear it out. Use these reports to update the AV signatures in your machine learning models.
- Use Attack Simulator to help educate people. Attack Simulator is a new feature offered in Office 365 Threat Intelligence. With Attack Simulator, you can send simulated phishing emails to anyone in your organization. These simulations help teach people how to handle an attack and give admins a better understanding of who is more susceptible to phishing. Attack Simulator is a powerful tool to test peoples' response to common phishing threats. It gives a better understanding about which people and/or groups might need more education or more rigorous protection policies.

Benefits

Better threat intelligence and cross-platform integration enhances individual services and makes it easier than ever for IT security pros to protect their people and companies against cybersecurity threats. AI and machine learning continue to improve, along with detection capabilities. Policies can be refined, configurations can be updated and there are with fewer infrastructure requirements. Using EOP and Office 365 ATP, we balance productivity and protection against advanced and sophisticated phishing campaigns.

Exchange administrators and security analysts in CSEO are saving time and responding faster to phishing at Microsoft. Security enhancements in Office 365 continue to give us best-in-class protection against the evolving threat landscape.

Improved visibility

With telemetry from EOP, Threat Explorer in Office 365 Threat Intelligence, and antivirus detection on malicious files, CSEO has more visibility of phishing attacks in our environment. Both ATP Reporting and Threat Explorer give us threat details and help us to triage an event. We can see who received the mail, who reported it, and who clicked a link or attachment. Threat Explorer makes it easier for us to spot trends when a phishing email is part of a targeted campaign against a specific user or role.

To investigate phishing emails before Office 365, we relied on employees to forward suspicious email to us. It required a lot of manual investigation to determine the nature and scope of the attack. It could take as long as eight days to investigate an attack—and often, we still would not have the full picture. With Threat Explorer, we can do in a single day what used to take more than a week. With these enhanced detection and response tools, we are seeing more than an 80 percent reduction in threat investigation times.

Self-service remediation

Using Threat Explorer, at CSEO we can search for and purge emails on our own—without having to rely on other teams. In the past, depending on the size of the phishing campaign, searching for malicious emails and engaging with the team that could purge them could take us days. Now we can search for malicious email, delete malicious attachments from mail, and/or move phishing email to the Junk folder.

By integrating technology platforms in our security stack, we have more detail about what happened before, during, and after an attack, and we can be agile in our efforts to protect our environment. For example, integrating Windows Defender ATP and Office 365 Threat Explorer now shows us who received the phishing mail, who opened it, and which client devices may have downloaded an attachment. Windows Defender ATP can quickly quarantine that email.

Strengthening our security posture

At CSEO, we have improved awareness, gained more insight, and increased productivity to address phishing. With the time and resource savings, we can be more proactive in strengthening our security posture against phishing campaigns. Since deploying the "Report phishing" functionality, we have seen a 37 percent increase in reported phishing and social engineering campaigns, and we have more capacity to handle those threats quickly. Our rich intelligence helps us identify trends, and we are more agile at updating AV signatures in our machine learning models. Our Exchange admins are also more agile and can quickly fine-tune phishing policies for people and groups across Microsoft.

For more information

Microsoft IT Showcase

microsoft.com/ITShowcase

Microsoft uses threat intelligence to protect, detect, and respond to threats

Set up Office 365 ATP Safe Links policies