axians

CYREN

Beyond the Email Perimeter

An analysis of hidden email inbox threats

Table of Contents

Introduction

The actions of cybercriminals have forced organizations to accept that their email perimeter defenses are incapable of preventing well-crafted email attacks. This is not to say those email perimeter defenses are unnecessary. However, while they are adept at blocking well-known phishing threats and viruses, they offer inadequate detection of social engineering email threats like business email compromise.

To close this gap, cybersecurity leaders have focused on employee security awareness training and processes for employees to report suspicious emails to the security operations center for analysis. However, relying on employees to spot latent email inbox threats does not provide leadership with complete visibility into the threat. What about the email attacks that go unreported? This approach also distracts employees from their primary roles and burdens security teams with a high volume of user-generated alerts.

In this report, the Cyren research team sheds light on the scale of social engineering attacks within the email perimeter to help cybersecurity leaders better understand the nature of the threat so they can optimize their abilities to hunt and remove malicious mails before a distracted user takes the bait. We also strongly recommend that businesses deploy complementary security technologies to help detect these threats and automatically manage them, making better use of employees' time in the process.

Key Findings

During an average month, there are 75 malicious messages per 100 mailboxes that slip past email security filters like Microsoft 365 Defender. The cost to respond to these threats is a serious concern. For example, a business with 5,000 mailboxes would need to detect and respond to 3,750 confirmed malicious inbox threats each month.

The per capita rate of malicious inbox content has more than doubled in the two years Cyren has been tracking this metric, even after we exclude content that borders on junk mail (e.g., spam/scam category in Figure 4). This suggests attackers are improving their tactics faster than email security filters are improving their detection.

The majority (79%) of these threats are phishing – emails containing URLs to web content intended to harvest login credentials, personal information, or payment details. Phishing has been the origin of many high-profile breaches and ransomware attacks; we and others suspect it is the most frequent precursor to more damaging attacks. Phishing may not be the hottest buzzword in the cybersecurity space, but it is clear to us that it remains a massive problem.

Specialized detection provides dramatic improvement of the "catch-rate" for inbox threats. The Cyren detection pipeline identified 91% of threats in the inbox at the time of arrival, 8% were flagged retroactively (probably due to delayed detonation), while the remaining 1% were reported by customer end-users.

Outsourcing incident response results in a significant reduction in Mean Time to Respond (MTTR) to email threats. The Cyren Incident Response service has a MTTR of 11 minutes compared to 38 minutes for organizations that rely on in-house teams to investigate suspicious email alerts. The most likely explanation for the difference in MTTRs is the massive volume of alerts of all types that in-house SOCs must manage.

Data and Information Sources

Cyren continuously detects malicious email, file, and web content on a global scale. This effort requires relentlessly collecting and automatically analyzing billions of suspicious objects each day.

The raw data for this report was extracted from business email traffic received by Cyren Inbox Security global customers from February through May of 2022. Cyren Inbox Security customers have a mix of third party and cloud-native email security.

The threat and incident response data were derived from the raw data by the Cyren Inbox Security detection pipeline and Cyren Incident Response analysts.

The Cyren Inbox Security detection pipeline continuously monitors email mailboxes for malicious content and abnormal activity. By analyzing email metadata and contents with sender and recipient behaviors it can detect threats that have not been blocked at the network perimeter by email security filters. The Cyren pipeline classifies message contents as clean, malicious, or suspicious. Suspicious content reported by the detection pipeline or end-users is analyzed by Cyren analysts for final classification.

Industries and Geographies

The industries represented in this research are listed below.

- Consumer Services
- Construction
- Education
- Energy
- Entertainment
- Finance
- Food and Beverage
- Government
- Healthcare
- Legal
- Manufacturing
- Not for Profit
- Real Estate
- Retail
- Technology
- Transportation
- Travel

Based on the location of company headquarters, the following geographies were included in this research. It is important to note that many customers have offices beyond just their headquarters address so this research likely applies globally.

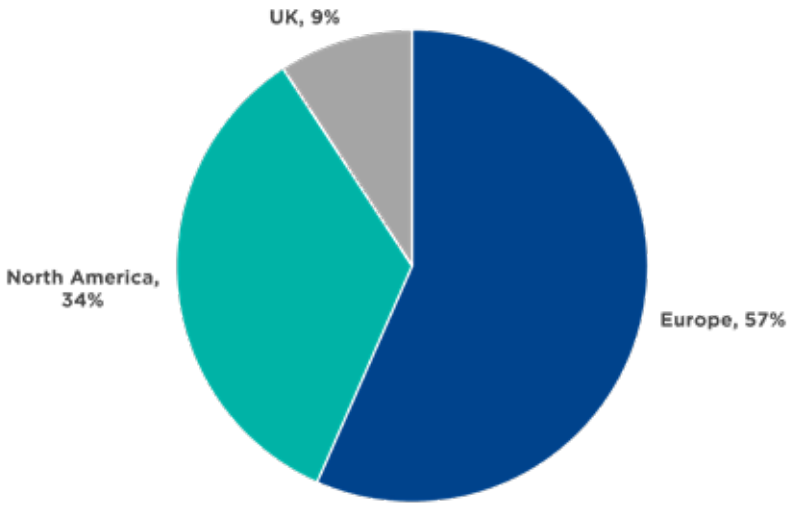

Figure 1
Distribution of malicious incidents by HQ region

Email Perimeter Security

As organizations move to the cloud, many migrate away from third party secure email gateways as their email security filter. The Cyren customer base reflects this trend. The chart below shows that almost three quarters of the organizations have adopted Microsoft 365 Defender as their primary email security platform.

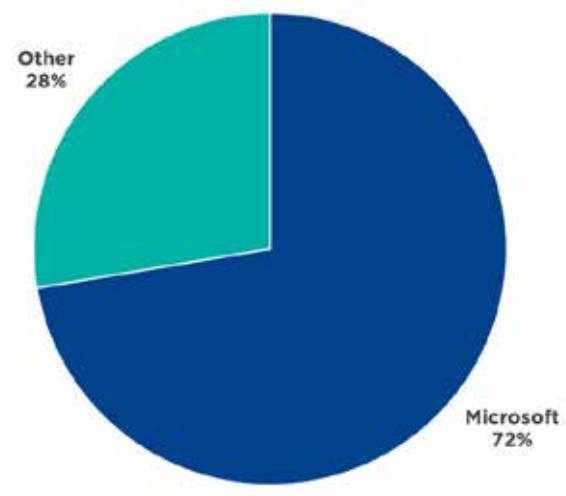

Figure 2
Native vs third party email security

The Detection Gap

Many organizations lack visibility of the total volume of threats affecting their users. They know how many emails were blocked their email security filter. Some may have accurate accounting of how many emails were flagged by employees. Neither of which provides a true picture of how many threats got through the email perimeter to users' mailboxes.

Rate of Mailbox Infection

During the time covered by our research, Cyren detected a monthly average of 75 malicious messages per 100 mailboxes. The infection rate has nearly doubled since Cyren began tracking it two years ago. Some of this increase can be attributed to detection improvements but the main contributing factors are that criminals have refined their abilities to slip past perimeter defences and are launching their attacks at scale.

75

Confirmed threats per 100 mailboxes

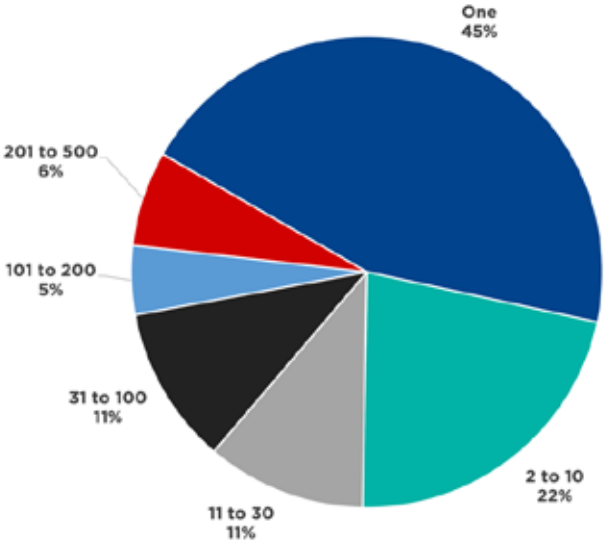

Figure 3
Number of recipients per attack

Targeted Attacks

More than half (55%) of all phishing, impostor, and malware email attacks detected were received by more than one user (Figure 3). However, two-thirds of email attacks were delivered to between 1 and 10 users suggesting the attacks were targeted. Targeted or unique email threats will be more difficult to block.

Threats Detected in Mailboxes

Accounting for 79% of malicious incidents (Figure 4), phishing continues to be the most frequent attack type. It is also the most likely precursor to account takeovers, convincing BEC attacks, and ransomware infections. Simply put, if users are exposed to fewer phishing threats, they are less likely to expose their login credentials which denies attackers the access required for the next step in their attack. More effective management of phishing threats including, specialized detection and automated incident response, will dramatically reduce the risks of all email attack types.

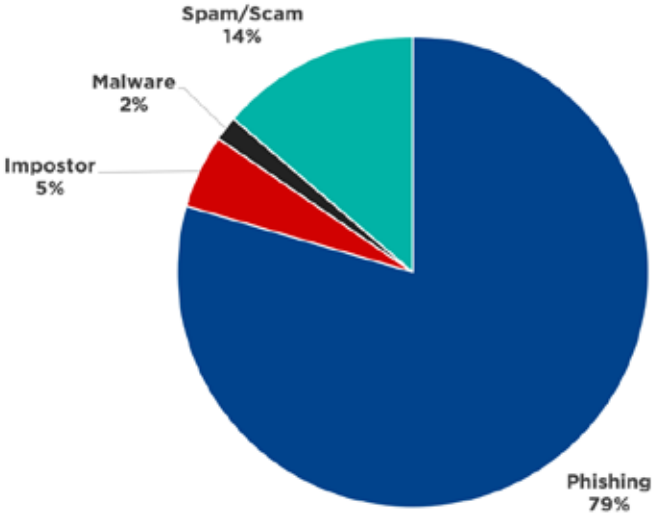

Figure 4
Malicious incidents by threat type

The prevalence of threat types varies by industry. For example, Construction industry customers experienced a much higher rate of phishing (Figure 5), while we detected a greater percentage of malware in email mailboxes of customers in the Manufacturing vertical (Figure 6). It's critical that businesses have this level of visibility to properly prioritize the management of risks like account takeover, ransomware, and financial fraud.

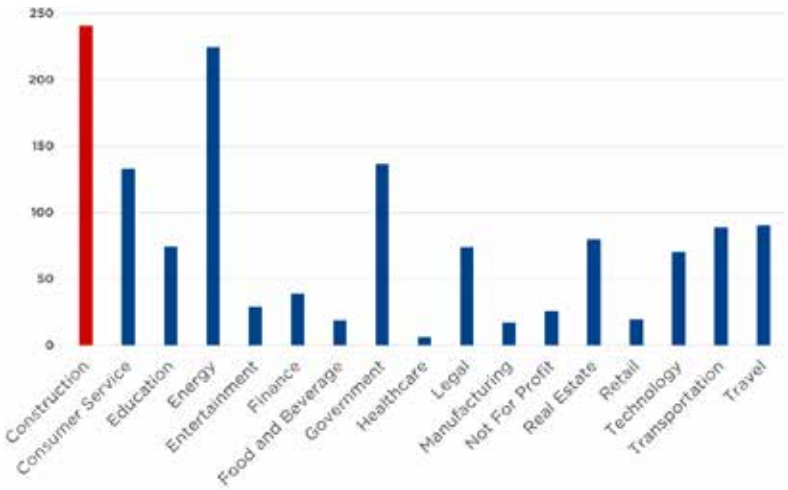

Figure 5
Monthly Phishing Incidents per 100 Mailboxes

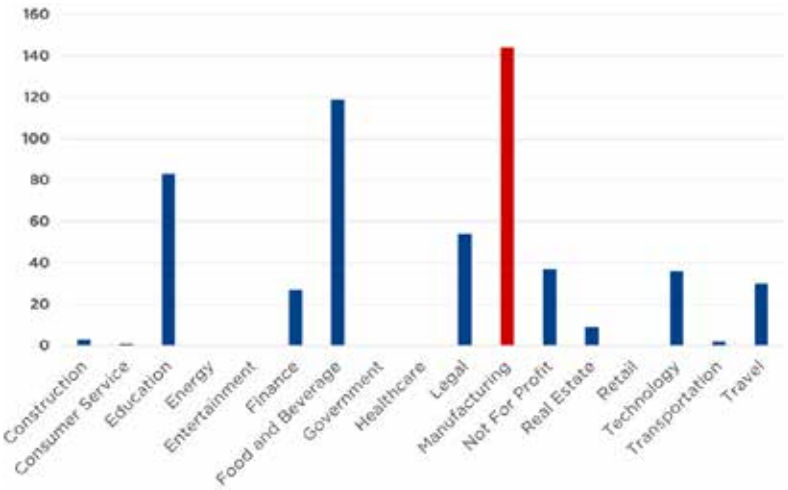

Figure 6
Monthly Malware Attachments per 100 mailboxes

We also looked at why these incidents that were classified as phishing, malware, etc. Credentials continue to be
the most common threat indicator detected by Cyren (Figure 7). These statistics reinforce that exposed login credentials are a prerequisite to many follow-on attacks. Obviously, better phishing defense is not the only solution to the account takeover problem. Multi-factor authentication (MFA) needs to be enforced for all systems that support it.

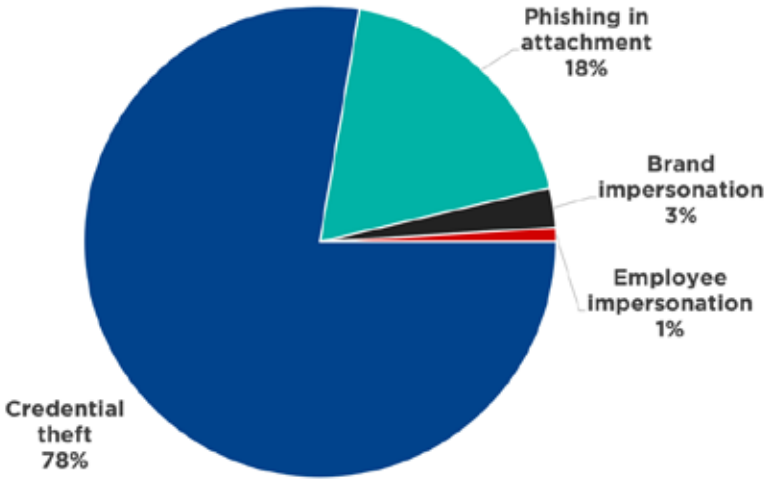

Figure 7
Phishing threat indicators

Tricks of the Trade: Evasion Techniques

Previous sections of this report include information to help cybersecurity leaders understand the scale of the threats evading detection and arriving in users' mailboxes. This section includes information about the techniques attackers use to evade detection by the email perimeter. Of course, this analysis is not exhaustive. While there may be a finite number of evasion techniques, attackers can combine them in an infinite number of ways.

Sender domain reputation and authentication are common ways that secure email gateways filter email messages. As Figure 8 shows, the most frequent technique attackers use to evade detection is sending emails from a well-known webmail domain like gmail.com.

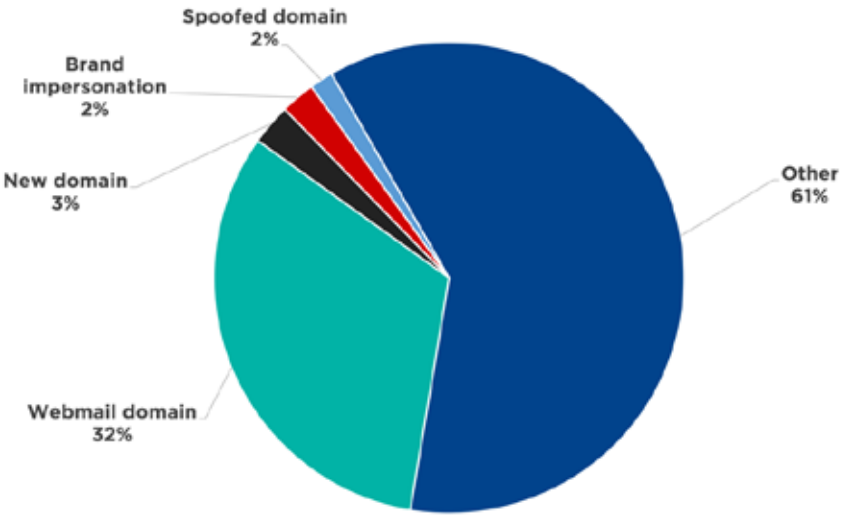

Figure 8
Sender domain evasion techniques (excludes Spam/Scam incidents)

Cybercriminals also must hide the intent of the email payload and destinations. They know that URLs and attachments get scanned by security engines and have perfected ways to avoid detection at the email perimeter. They also play a cat-and-mouse game with the cloud companies whose services they abuse to host phishing sites, malware, and other malicious content. Therefore, businesses need to close the detection gap by continuously scanning mailboxes for threats and optimizing their incident response processes to remove those threats as fast as possible. Figure 9 shows how attackers evade detection of their phishing pages and hosted malware. Note the use of new domains for both sending emails (Figure 8) and hosting malicious content (Figure 9).

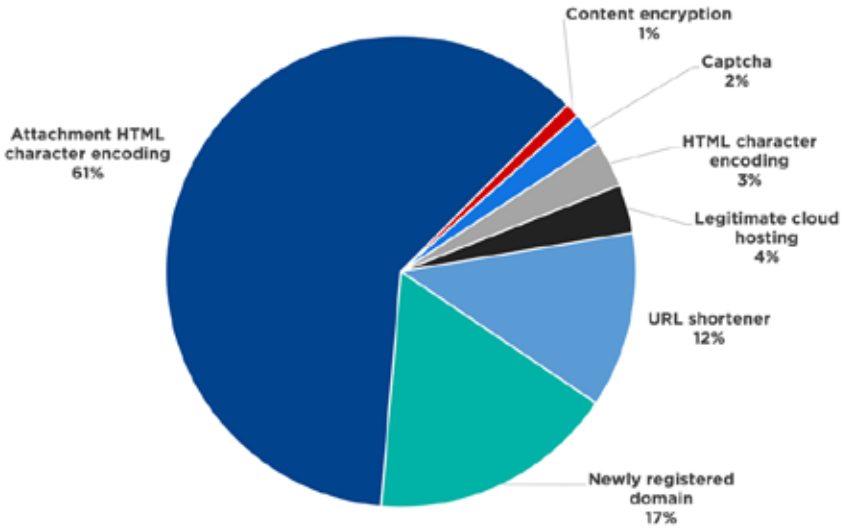

Figure 9
Content and hosting evasion techniques (excludes Spam/Scam incidents)

Efficacy of Detection Models

Detecting inbox threats requires a variety of approaches including real-time analysis, threat intelligence, and user-reported. We define each approach as follows.

- **Real-time analysis** detects inbox threats when they are active at the time of initial inspection. Techniques include natural language processing, real-time content analysis, and user entity behavior analytics to classify message contents as clean, malicious, or suspicious when the message arrives.
- **Threat intelligence** detects threats that become active after the initial real-time analysis. It requires keeping a record of message metadata like URLs, file attachment fingerprints (file hashes), sender addresses, etc. and continuously comparing them to the constantly evolving universe of malicious objects to retroactively reclassify a clean message as malicious or suspicious.
- **User-reported** detection is a manual model that relies on end-users to spot suspicious messages and submit them to security analysts to investigate the messages.

No single model can detect 100% of inbox threats. However, automated methods like real-time analysis and threat intelligence play an oversized role compared to user-reported (manual) detection (Figure 10).

One tactic that cybercriminals use to evade detection is delayed detonation. The aim of delayed detonation is to avoid real-time detection by de-activating threats during the time of inspection. The easiest example of this is an email that contains innocent language and a URL to a trusted website with no malicious content. The criminal sends the email and waits 15 minutes before redirecting the URL to a different, malicious web page. As shown in Figure 11, the Threat Intelligence detection model helps to identify delayed detonation threats by retroactively classifying emails based on new information and updates to detection logic.

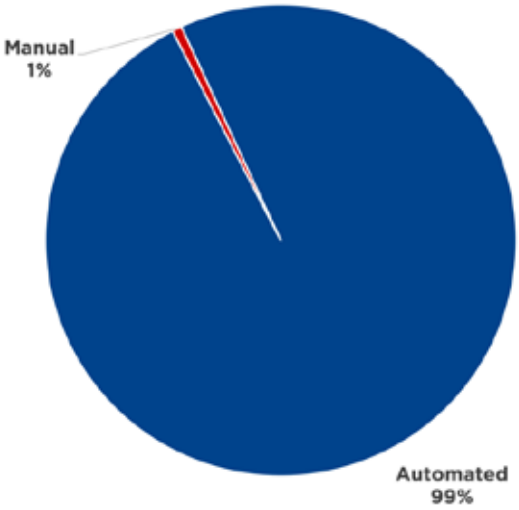

Figure 10
The nature of threat detection

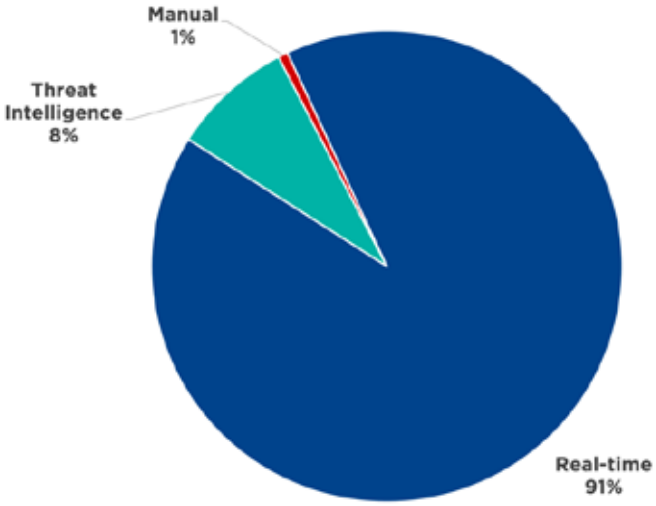

Figure 11
Delayed detonation of threats

Manual Detection in Context

Cyren researchers found that automated detection models spotted 99% of the confirmed threats. The 1% that required manual analysis were messages classified as suspicious or clean. Classified as suspicious meaning the automated detection models identified something out of the norm associated with the content but could not with high confidence determine if the message was harmless or a threat. Classified as clean meaning a false negative.

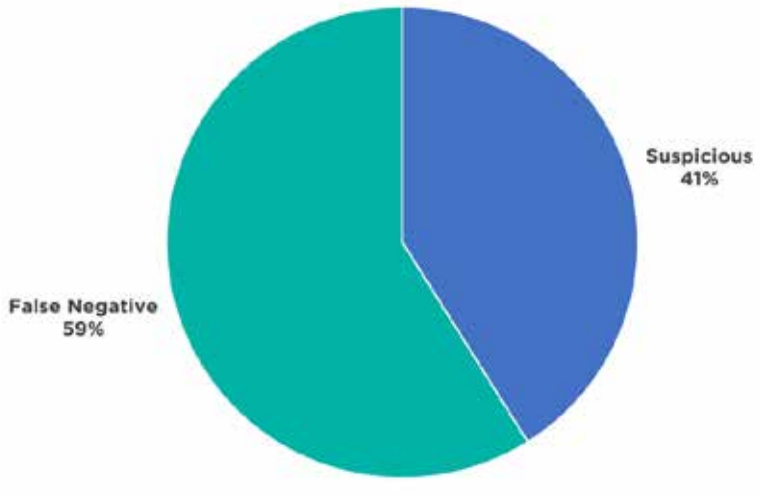

Figure 12
Origin of incidents that required manual analysis

Bear in mind the chart to the left represents 1% of the total malicious incidents detected by Cyren during the period covered by this research. Even though this represents a small number of threats, it does not diminish the importance of establishing a culture of security and optimising incident response.

Classifying messages as suspicious and displaying the indicators to users is a recent addition to the defense against email threats. This is a massive improvement over static banners like “This email originated outside the organization” which occasionally help users spot impostor emails. This approach, shown in Figure 13, strengthens the knowledge imparted by user security awareness training by enabling users to apply lessons learned in real-time and to (potentially) real threats.

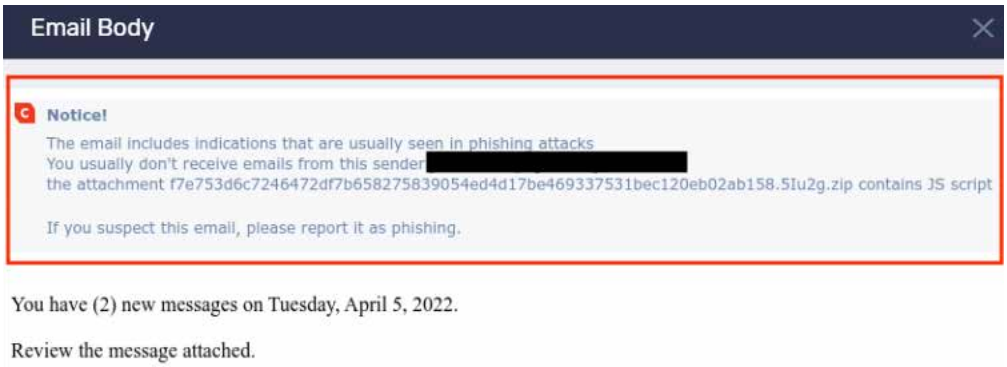

Figure 13
Suspicious indicators displayed to the user

Mean Time to Respond

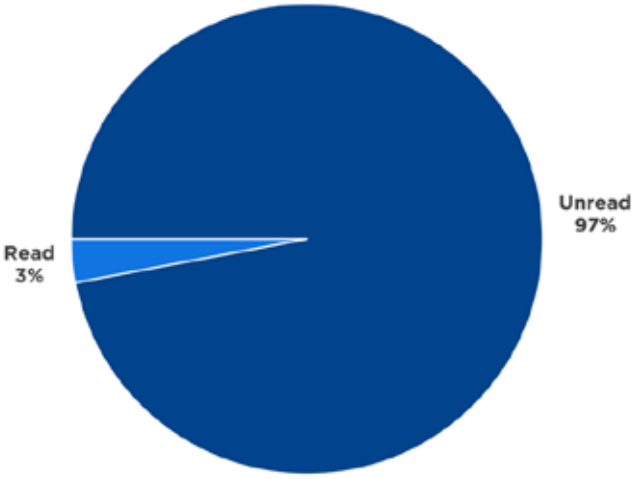

Figure 14
Malicious message state at time of remediation

Mean Time to Respond is the average elapsed time from when a malicious email is reported to when the threat is neutralised. When threats are both detected and eliminated automatically, the MTTR is a matter of seconds. The short MTTR of automated incident response processes is implied by the high percentage of malicious messages that were still marked as unread by the user.

Threats that cannot be automatically detected, cannot be automatically removed. Incident response teams must investigate the messages, then locate every instance of the threats and remove them before they can cause harm. Outsourcing that investigation and response yields a MTTR that is less than a third of the time it takes when organisations add that responsibility to their existing teams (see Figure 15). Note, the MTTR of in-house teams is likely to be much higher for organisations to do not have the benefit of responding to only 1% of malicious inbox content.

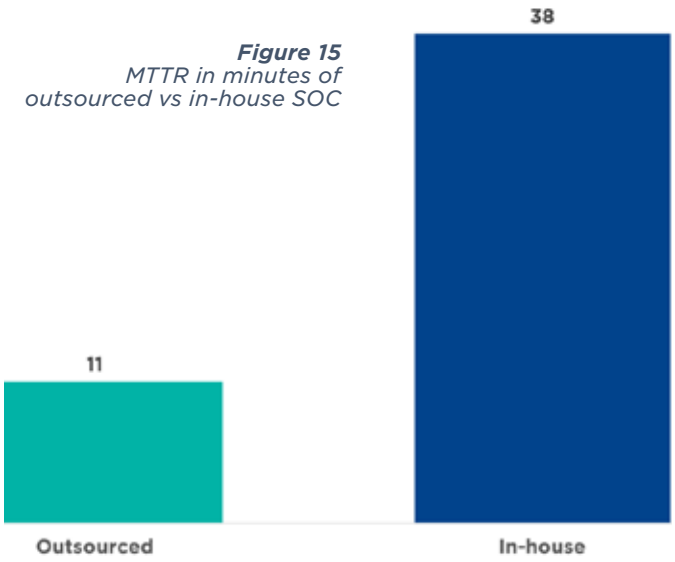

Figure 15
MTTR in minutes of outsourced vs in-house SOC

Conclusion

The scale of threats inside the email perimeter is high at a monthly average of 75 confirmed threats per 100 mailboxes. Organizations with a lower rate either do not have a complete picture of the problem or are less of a target than the sample used for this research. It should not be a question left unanswered.

Phishing continues to be the dominate threat type, representing 79% of malicious content detected in mailboxes. While threats like ransomware, account takeover, and vendor email compromise may be “on trend,” this data suggests organizations can greatly improve their ability to pre-empt those issues by optimizing their abilities to detect and quickly contain evasive phishing. Of course, threats like ransomware and account takeover can occur without exploiting the email attack vector. Best practice password policies, multi-factor authentication, endpoint security software, and other common guidance must be followed.

Businesses have spent billions on secure email gateways and other filters designed to block threats from delivery, but the problem of social engineering emails persists. The typical response to this persistence of mailbox threats has been to train users to spot, avoid, and report email threats. However, this approach exhausts the most precious resources within an organization: its people. Better to deploy layered detection models, automated remediation, and outsourced incident response to complement existing investments in email perimeter defense, user training, and security operations. This approach will reduce the current costs of manual detection models and existing incident response workflows and yield a far better MTTR.

About Cyren

Cyren protects more than a billion users around the world from sophisticated and emerging email-, malware-, and web-based cyber-attacks every day. Our embedded threat detection, threat intelligence and inbox security solutions help enterprise, service providers, and technology companies prevent breaches and eliminate countless hours of incident response.

Learn more about our solutions at www.cyren.com.

Authored By

Ira Chernous
Lior Kohavi
Mike Fleck

Cyren is a messaging security company that protects enterprise email users from today's evasive threats and supplies threat intelligence solutions to security software integrators, hardware OEMs, and large service providers. Cyren's GlobalViewTM threat intelligence network analyzes billions of email and web transactions daily and is trusted by companies like Microsoft, Google and Check Point, who utilize Cyren's APIs and SDKs to operationalize threat intelligence for their customers.

HEADQUARTERS

US Virginia
1430 Spring Hill Road Suite 330
McLean, Virginia 22102
Tel: 703-760-3320
Fax: 703-760-3321

SALES & MARKETING

UK Bracknell
Maxis 1
43 Western Road
Bracknell Berkshire RG12 1RF

US Silicon Valley
1230 Midas Way Suite 110
Sunnyvale, CA 94085
Tel: 650-864-2000
Fax: 650-864-2002

R&D LABS

Germany
Hardenbergplatz 2
10623 Berlin
Tel: +49 (30) 52 00 56 - 0
Fax: +49 (30) 52 00 56 - 299

Iceland
Dalshraun 3
IS-220, Hafnarfjordur
Tel: +354-540-740

Israel
1 Sapir Rd. 5th Floor, Beit
Ampa P.O. Box 4014
Herzliya, 46140
Tel: +972-9-8636 888
Fax: +972-9-8948 214

Cyren.com

@CyrenInc

linkedin.com/company/cyren