E-BOOK

Turning the Tables

How smart email reporting and remediation can transform attackers' favourite targets into your best defence

proofpoint.

Introduction

Even before this decade's pandemic-fueled migration to remote and hybrid work, cyber criminals had shifted their focus away from infrastructure and toward people. More than ever, attacks seek to exploit human vulnerabilities, not just technical flaws. In most cases, they do it through email—a still-ubiquitous communications platform that was never built with security in mind.

According to the *2021 Verizon Data Breach Investigations Report*, 85% of breaches involve "the human element,"[1] with email as the preferred conduit for social engineering.[2]

$15 million

The **cost of phishing** has nearly quadrupled to almost $15 million per year for organisations, or $1,500 per employee.[3]

2X

Phishing attacks doubled in 2020 alone.[4]

3/4

About three-quarters of all ransomware originates from email phishing, the top attack vector for this growing threat.[5]

1 Verizon. "Data Breach Investigations Report Executive Summary." May 2021.
2 Verizon. "Data Breach Investigations Report." May 2021.
3 Ponemon. "The 2021 Cost of Phishing Study." August 2021.
4 APWG. "Phishing Activity Trends Report 4th Quarter 2020." February 2021.
5 Unit 42, Palo Alto Networks. "Ransomware Families: 2021 Data to Supplement the Unit 42 Ransomware Threat Report." July 2021.

Email is universal, critical to modern business and inherently insecure. Created long before the internet was mainstream, email was never developed with privacy or security in mind. In the 45 years since, it has become an essential pillar of modern business communications—and a magnet for all kinds of attacks.

Every day, billions of emails reach users' inboxes. Most are benign. Many are nuisances. Some are important. And all too many are dangerous. Fully protecting your organisation against these threats requires more than one layer of defence.

The good news: you can use attackers' tactics against them. By making email reporting and remediation key parts of a multilayered defence, you can transform every potential victim into a defensive choke point.

This e-book explains how to teach users to recognise and report suspicious emails—without creating needless work by forcing IT and security teams to chase down false alarms.

It includes new research into user behaviour around simulated email threats. It also explores ways of getting more users to report suspicious emails. And it maps out practical steps you can take to streamline your response to suspicious email reports while helping users get better at discerning real threats.

SECTION 1

User Reporting Rates in Phishing Simulations

Phishing simulations in security awareness programmes have always been a popular tool for understanding user behaviour. But the “click rate” or “failure rate” of simulations is not an ideal measure of user behaviour when it’s the only metric considered.

Making email reporting easier—for users, IT and security teams

A convenient way for users to report phishing is through an abuse mailbox with an email address—for example, phishing@abccompany.net. While abuse mailboxes are effective, they often require back-and-forth between IT and users to gather critical details like email headers.

That's why email reporting add-ins or buttons, like Proofpoint PhishAlarm, which feature simplicity and more functionality than an abuse mailbox, are becoming popular tools for reporting suspicious email messages.

The failure rate is defined as the worst-case scenario based on the type of simulated phishing email. In short, did the user fall for the bait and respond by clicking the link, opening the attachment or entering credentials or personal information?

You can make reporting even more accurate by including context that helps users identify malicious emails. For example, the "Report Suspicious" button on our email warning tags feature can alert users to emails that employ common phishing techniques such as spoofing and lookalike domains.

When deploying such a tool, look for solutions with HTML-based banners that are contextual, customisable and work with little to no IT overhead.

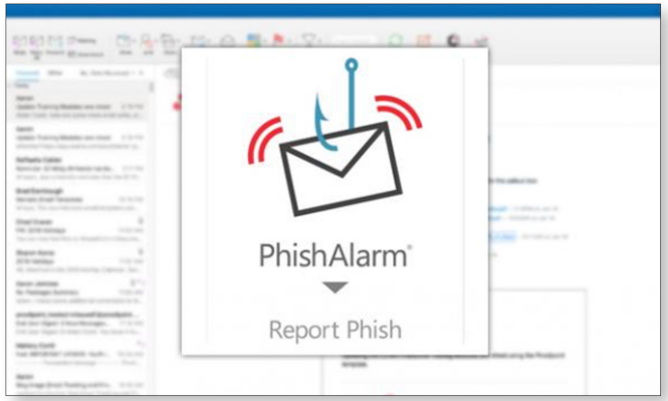

Figure 1: Proofpoint PhishAlarm email reporting button.

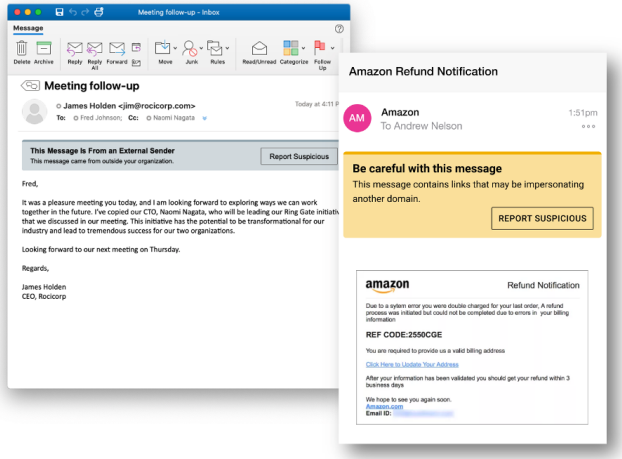

Figure 2: Proofpoint email warning tags with a "Report Suspicious" function will help customers improve their email reporting accuracy and email security.

When failure isn't enough

But the failure rate alone isn't the best way to measure user behaviour. First, the failure rate can vary widely according to the type of simulated phishing email used, making it hard to see patterns or trends.

For example, some less challenging phishing templates have low single-digit failure rates. Others, such as the Netflix template in Figure 3, have netted failure rates approaching 100% in certain campaigns.

Getting a fuller picture

When it comes to understanding user risk, failure rates give you just a small part of the picture. For a more complete view, you also need to take into account the reporting rate of phishing simulations.

The reporting rate demonstrates that users are not just "avoiding the bad" but also "doing the good" and helping to secure the organisation. The higher your reporting rate, the more likely users are to report actual suspicious messages to your IT and security teams, helping them resolve true threats before they cause lasting harm.

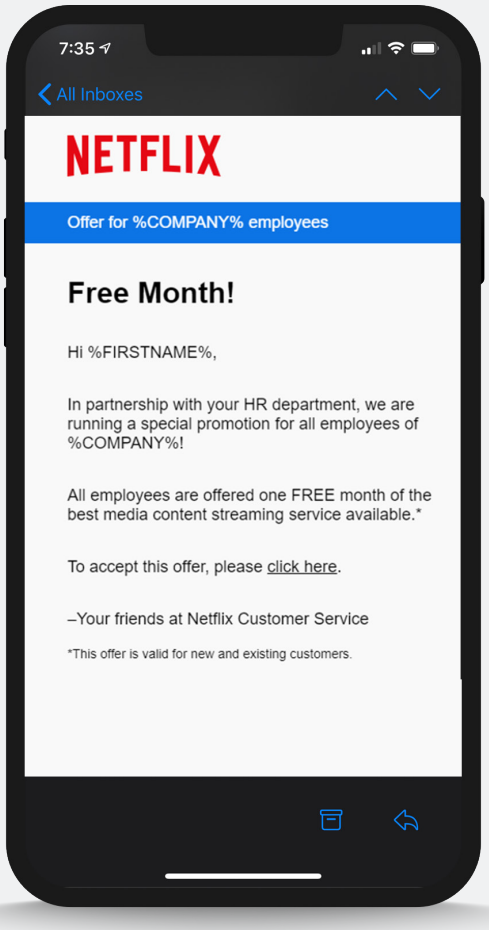

Figure 3: Template for simulated phishing email using a fake Netflix promotion as a lure.

Figure 4: Using this formula, the average resilience factor for Proofpoint customers is 1.2.

Scoring your resilience factor

When you combine users' failure rate with their reporting rate, you get what we call the "resilience factor."

How this formula works:

1. **Start with the average reporting rate for phishing simulations.** For the average Proofpoint customer, that's 13%. (The average number of messages reported per user was just over five per year. We expect this figure to grow strongly in the coming years as organisations ramp up their education initiatives and users get comfortable with reporting suspicious messages.)
2. **Divide that number by the average failure rate for those simulations.** For the average Proofpoint customer, that's 11%. Your failure rate will almost never be 0% and can vary drastically depending on whether you make simulations more targeted and difficult.
3. **The result is your resilience factor.** For Proofpoint customers, it's 1.18, rounded up to 1.2 in Figure 4.[6] We recommend a resilience factor of 14 (an average reporting rate of 70% and a failure rate of 5% or less) as a stretch goal.

Your resilience factor is one of the most reliable indicators of user risk. Some customers have achieved a factor of 14. With the right training programme and a culture of security awareness, you can meet and even exceed that level. But changing user behaviour takes time and constant engagement.

6 Keep in mind that we serve customers at every stage of their security awareness journey. This average includes resilience factors for programs that are just starting out and those with mature programs.

14

Your resilience factor is one of the most reliable indicators of user risk. With the right training programme and a culture of security awareness, some customers have achieved a resilience factor of 14.

Reporting and failure rates by type

Call it the attachment paradox: attachment-style simulated phishing emails drive the highest reporting rate—and the highest failure rate (see Table 1). Users may find attachments tempting because of the data promised in them (such as COVID-19 exposure data, employee bonus numbers and the like). But a certain segment of users may be more (rightly) cautious about attachments and report them at a higher rate.

SIMULATED PHISHING TYPE	REPORTING RATE AVERAGE	FAILURE RATE AVERAGE	RESILIENCE RATIO AVERAGE
ATTACHMENT	18%	20%	0.9
DATA ENTRY/CREDENTIAL	15%	4%	3.8
LINK	13%	12%	1.1

Table 1: Failure and reporting rate, by type.[7]

The resilience ratio for data entry/credential phishing simulations was much higher than the other two types. But this type of attack requires an additional step from users—to fail, they must click on the link and submit their credentials.

7 The reporting and failure rate figures in this section are from our *2021 State of the Phish Report*.

A closer look at resilience factors by industry

Table 2 shows an overview of industry-specific data on failure and click rates along with resilience factors. Financial services had the highest reporting rate (20%). But the legal industry has the highest overall resilience factor (2.1).

INDUSTRY	REPORTING RATE	FAILURE RATE	RESILIENCE FACTOR
FINANCIAL SERVICES	20%	11%	1.8
ENERGY/UTILITIES	18%	11%	1.6
INSURANCE	17%	10%	1.7
LEGAL	17%	8%	2.1
ENGINEERING	16%	16%	1.0
AUTOMOTIVE	15%	8%	1.9
BUSINESS SERVICES	14%	11%	1.3
TECHNOLOGY	13%	12%	1.1
GOVERNMENT	13%	10%	1.3
MINING	13%	13%	1.0
FOOD & BEVERAGE	11%	11%	1.0
MANUFACTURING	10%	10%	1.0
HEALTHCARE	10%	10%	1.0
ENTERTAINMENT/MEDIA	10%	9%	1.1
TRANSPORTATION	10%	12%	-1.2
TELECOMMUNICATIONS	9%	14%	-1.6
CONSTRUCTION	9%	11%	-1.2
RETAIL	9%	13%	-1.4
EDUCATION	6%	12%	-2.0
HOSPITALITY/LEISURE	5%	10%	-2.0

Table 2: Reporting and failure rates, by industry.

Improving user reporting rates

As seen in Figure 5, a small number of organisations in our data set account for the largest number of reported simulated phishing emails. Most organisations fall into a low range.

Reporting rate performance depends on the organisation's security awareness maturity and how long users have been using an email reporting add-in or button. Often, low performance may indicate low awareness of the reporting tool rather than the phishing threats themselves.

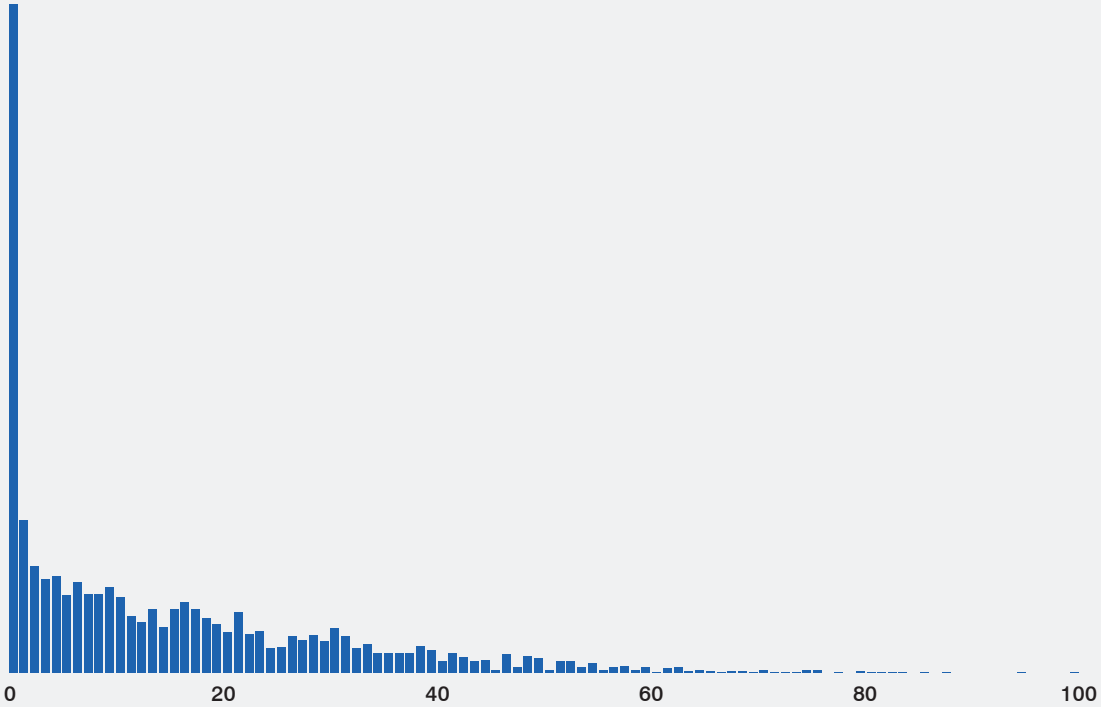

Figure 5: Distribution of reporting rates.

PERCENTILE	PERCENTAGE OF USERS REPORTING SIMULATIONS
25%	1.4%
50%	8.5%
75%	19.9%
Average	13%
TOP-PERFORMING	83.6%

Table 3: Average reporting rates, quartile breakdown.

With low reporting rates, we've often found that:

- Organisations have only recently deployed their email reporting add-ins
- Users may have multiple options to report messages (such as an abuse mailbox address) that are not reflected in our data
- User education does not include instructions for how to use an email reporting add-in
- Users instinctively delete or ignore messages they aren't expecting

A little education can go a long way toward helping your organisation raise the reporting rate. Explain how to use the reporting add-in as part of your security awareness training. Remind users of the tool in regular communications. And give feedback when users report simulated phishing messages. These steps can improve your reporting rate—and get you closer to being a top performer.

SECTION 2

Getting Real: User Reporting Rates in Actual Attacks

Measuring the reporting rate of phishing simulations is essential. But does it translate to real-world results? When a truly malicious or suspicious message reaches the inbox, do users recognise and report it? And how accurate are those reports?

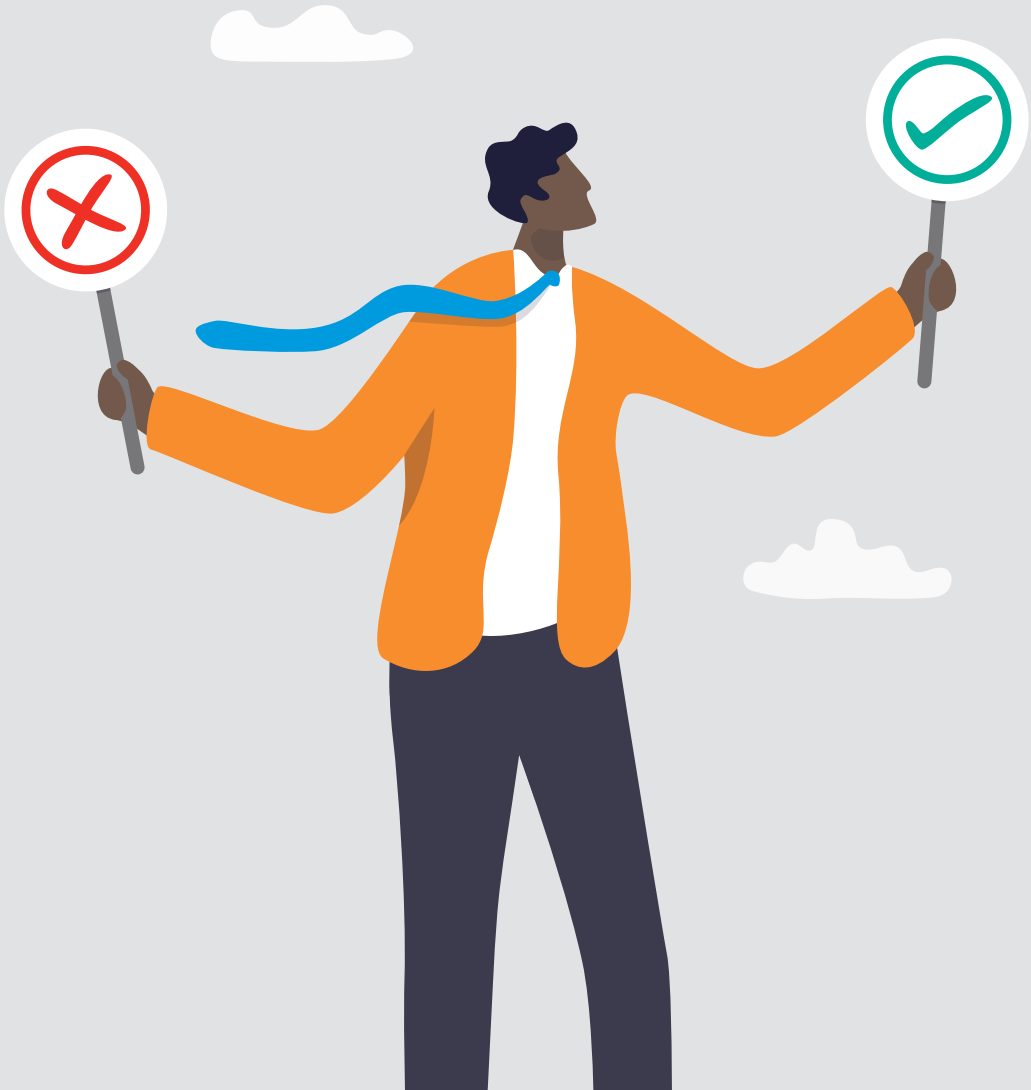

2.2 billion

Number of emails we analyse every day.

We set out to answer those questions by analysing data from the millions of people who reported messages they deemed suspicious using the Proofpoint PhishAlarm phishing button.[8] We gauged users' accuracy by comparing their reports with how those emails were classified by our email detection engine, which powers our Threat Protection Platform.

Note: Some might question how emails tagged as malicious by our detection engine would reach users' inboxes in the first place. Not every PhishAlarm customer uses our email security products. Often, we know the email is malicious because the attack has emerged somewhere in the billions of emails we analyse every day.

Detection categories include:

- **Malicious.** These emails contain malware, phishing, impostor emails or other threats.
- **Suspicious.** These emails are likely to be malicious, so you should quarantine them. But review them to ensure no legitimate email is lost.
- **Spam.** These emails are a nuisance and could contain malicious content.
- **Bulk.** These emails are the low-priority or promotional type. They don't pose a threat.
- **Low risk.** Closer analysis of these emails finds no signs of malicious content.
- **Unlikely a threat.** These emails don't reveal any malicious content or activity when analysed in a sandbox and by the detection stack.

We divided organisations into two distinct groups. The first group uses Proofpoint tools to report suspicious and malicious emails, but not spam. The second group uses Proofpoint tools to report spam along with suspicious and malicious email. Some organisations consider spam a "good" report. Others do not, which is why they were divided into these two groups.

8 Among suspicious emails reported between September 2019 and October 2020 through our PhishAlarm feature.

Malicious and suspicious messages

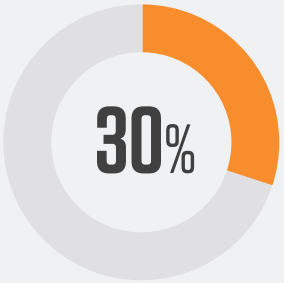

For most organisations, around 30% of reported messages are, in fact, malicious or suspicious.

Top performers can reach a number well above 50%.

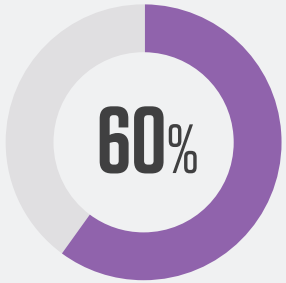

Many organisations achieve accuracies above 60%.

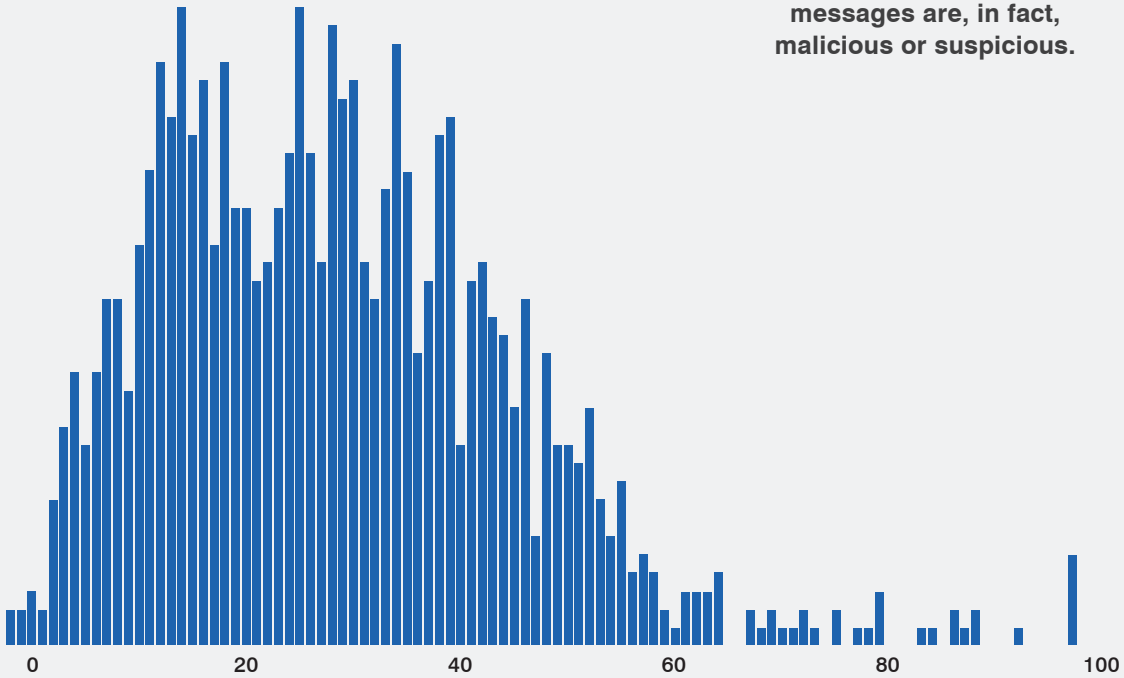

Figure 6: Accuracy of reported emails distribution.

PERCENTILE	ACCURACY
25%	18.1%
50%	29.6%
75%	41.1%
Average	31.0%
TOP-PERFORMING	100%

Table 4: Accuracy of reported malicious, and suspicious emails, quartile breakdown.

Malicious, suspicious and spam messages

When spam messages are included in the reporting, accuracy increased across the board, averaging about one-third. Many users may be unsure about what constitutes a malicious email, but most people know spam when they see it.

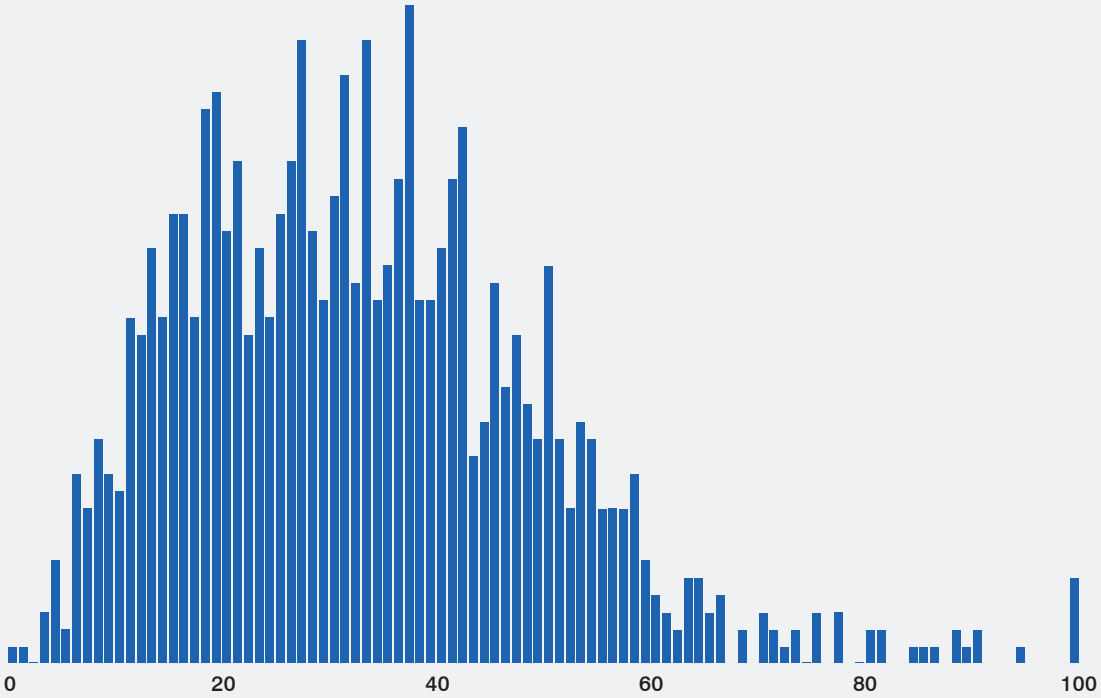

Figure 7: Email reporting accuracy distribution, spam included.

When spam messages are included in the reporting, accuracy averages about one-third.

PERCENTILE	ACCURACY
25%	20.4%
50%	31.7%
75%	42.6%
Average	33.1%
TOP-PERFORMING	100%

Table 5: Accuracy of reported malicious suspicious and spam emails, quartile breakdown.

SECTION 3

Breaking Loose from the Abuse Mailbox

The abuse mailbox can be a huge pain point for information security teams. It can mean thousands of hours spent:

- Researching messages
- Identifying which ones are malicious
- Trying to remove all copies before threats are potentially activated by users

And that's for real threats. For inaccurate reports—more than two-thirds of all reported messages for the average organisation—the cost is even greater. Already-stretched security and IT teams must chase down false positives while potentially serious attacks wait in the wings.

$700,000

Responding to credential phishing alone, one of the most common threats, costs organisations about $700,000 per year.[9]

The average 10,000-person organisation spends thousands of hours remediating threats such as business email compromise (BEC), malware infections, ransomware and credential theft, typically the main drivers of email-based attacks. Responding to credential phishing alone, one of the most common threats, costs organisations about $700,000 a year.[10]

TASKS	MALWARE INFECTIONS	BUSINESS EMAIL COMPROMISE	RANSOMWARE	CREDENTIAL THEFT
PLANNING	1,248	1,019	967	885
CAPTURING INTELLIGENCE	4,892	4,450	3,889	3,630
EVALUATING INTELLIGENCE	4,282	5,001	4,200	5,411
INVESTIGATING	12,045	12,336	11,901	12,884
CLEANING & FIXING	12,215	14,395	13,415	11,950
DOCUMENTING	951	1,075	913	1,002
TOTAL HOURS	36,633	38,276	35,285	35,762

Table 6. Hours IT teams spent resolving different types of phishing attacks. (Source: *The 2021 Ponemon Cost of Phishing Study*)

9 Ponemon. "The 2021 Cost of Phishing Study." August 2021.
10 Ibid.

SECTION 4

Next Steps: Make Reporting Easier, More Manageable

A security leader who was testing our PhishAlarm email reporting button once told us that while he liked the idea of user-enabled reporting, he wasn't ready to fully deploy it. His team was worried they'd be overburdened by user-reported false positives.

He had a point. On average, about two-thirds of the email messages that users report aren't actually harmful. That leaves a lot of room for improvement.

Fortunately, you can encourage user reporting and reduce risk without drowning in a sea of false alarms. **This section explains how.**

01 Start with what's in your inbox—the good, bad and ugly

If you're already overburdened with an influx of spam, phishing and other messages flooding your abuse mailbox, it's important to go to the very core of the problem: **messages getting delivered** to your users in the first place. That's why you need an **advanced email security** solution to reduce the volume of malicious messages in your inbox.

While we highly recommend an automated email reporting and remediation solution such as **Proofpoint Closed-Loop Email Analysis and Response (CLEAR)**, automation alone takes you only so far. At some point, the volume of bad email getting through—and then being reported—is too much for even the best automated systems. Fewer malicious emails reaching inboxes means dealing with fewer reports from users.

68%

In one case, a CESS claimed 900 messages were malicious. In reality, 68% were false positives.

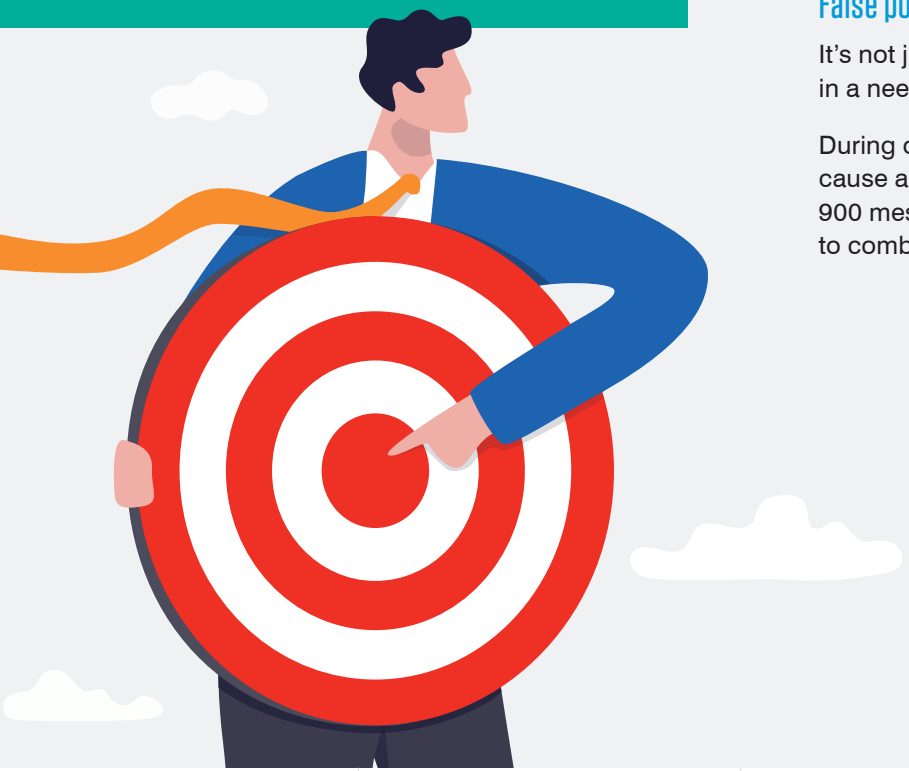

An ounce of prevention

Think of a healthy email inbox as preventive care. It's better to tackle the problem when it emerges rather than deal with the potentially worse downstream impacts later. This approach echoes frameworks such the MITRE ATT&CK "Shift to the Left" toward PRE-ATT&CK.

We frequently see this problem in email security when conducting our threat assessments with potential customers. In a recent set of proof-of-concept engagements we conducted with organisations using Microsoft email gateways, **we found hundreds of thousands of serious threats that had slipped through**. They included credential theft, malicious attachments, unsafe URLs and BEC threats.

The net result is more incident response time and resources, creating a downstream impact.

False positives are a big negative

It's not just actual threats and spam that cause pain for security teams. False positives also play a role in a needlessly large incident response workload.

During our proofs of concept, we've seen many API-based cloud email security supplements (CESS) cause a high false-positive rate for prospects during their proofs of concept. In one case, a CESS claimed 900 messages were malicious. In reality, 68% were false positives. That means the team would have had to comb through more than 600 inaccurately flagged messages manually to ensure they were delivered.

02 Help users report more accurately

Improving user email reporting accuracy can do wonders to reduce the number of false-positive reports. (That said, we still suggest that users report anything they're unsure about. It's better for your team to have visibility into any potential threat lurking in users' inboxes.)

Onboarding and ongoing reinforcement

Complement your phishing education initiatives by showing users how they can put their newfound skills to good use. You'll improve user reporting accuracy by explaining the email reporting button, what it's for, when to use it, and how to access it on different devices.

Here are a few ways to bring the reporting add-in to users' attention:

- Show the interface in company newsletters and explain how it works
- Discuss it during company town halls and in all-hands meetings
- Incorporate reporting information into your existing security awareness programmes

You can improve user reporting accuracy with a launch plan that works hand in hand with a comprehensive phishing awareness training programme. Use a mix of media and formats to engage users.

Feedback on reported messages

No one wants their efforts at vigilance to disappear into an organisational black hole. If users don't believe anyone is taking their reports seriously, they'll stop trying. A broken feedback loop can also mean new tickets for your team to manage if users follow up on their original report.

It's much easier to simply automate the feedback loop, letting users know the status of the message they reported. This feedback will sharpen users' skills and improve their reporting accuracy.

Nudge users with in-email warning labels

HTML-based email warning tags provide users with contextual nudges to enhance their email reporting volume and accuracy. Giving users context about the message—within the message itself—can help users gauge the risk in real time.

These nudges can be customised for different types of potentially malicious messages to get users' attention and allow them to report phish more easily.

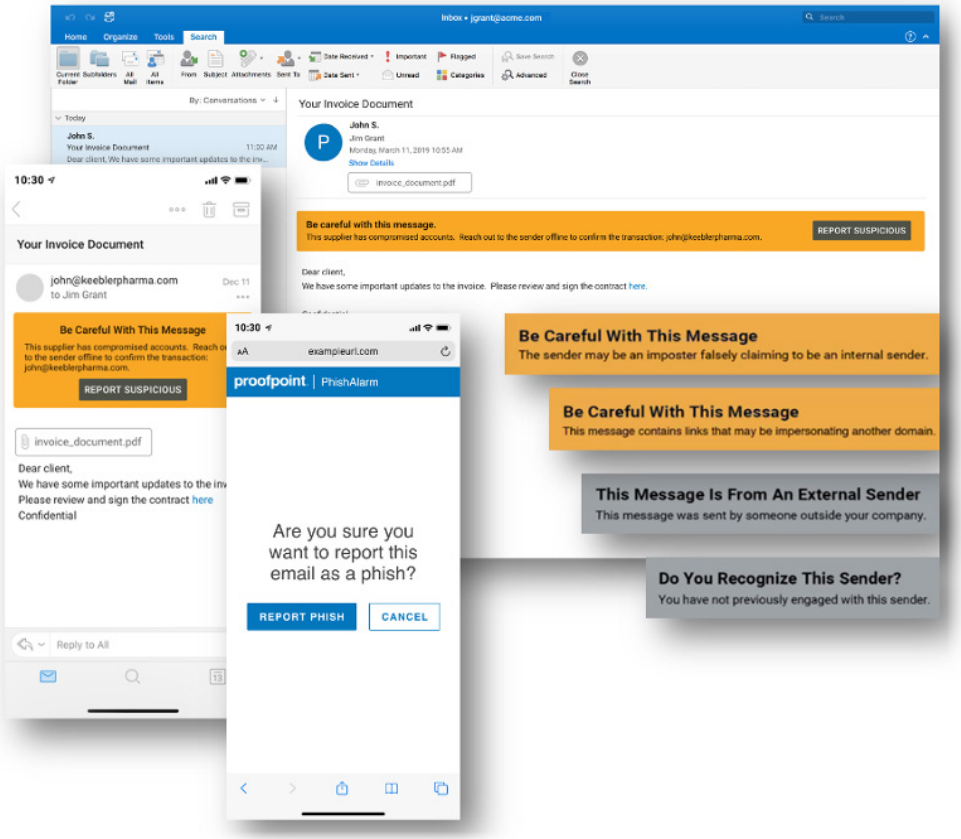

Figure 8: In-email warning labels can help direct users' attention to messages that may pose a higher risk.

03 Use automation and consider professional services

So you've improved your overall security posture and given users tools and guidance to improve their phishing email reporting. The last step is to automate as much of the incident response process as you can.

By automating the abuse mailbox, you can streamline workflows and slash the time spent on manual tasks. Some of our customers have reduced their workload as much as 90%.

But organisations are also strained for resources to manage email security. If that's a pain point for you, consider a best-of-breed **managed email service** partner to reduce IT burden and improve your email security posture.

Not every security team has the expertise or resources to piece together threat intelligence and sandboxing data. Fewer still have the staff or time to constantly update YARA rules and playbooks to stay on the leading edge. (YARA is a pattern-matching tool that helps with malware research.)

Professional services can help fill the gap, freeing up your team to focus on more strategic security tasks.

Conclusion

User-driven email reporting can be a critical part of your security strategy. But without the right approach to remediation, it can be a double-edged sword.

To avoid overwhelming your security team with email reports—and a flood of false positives—you must reduce the volume of malicious email spam, improve users' reporting accuracy and automate (or outsource) your response.

By following the strategies outlined in this e-book, you can improve your organisation's security posture and create a more streamlined abuse mailbox.

Learn how Proofpoint can help you build layered defences against phishing.
Visit **proofpoint.com**.

LEARN MORE

For more information, visit **proofpoint.com**.

ABOUT PROOFPOINT

Proofpoint, Inc. is a leading cybersecurity and compliance company that protects organisations' greatest assets and biggest risks: their people. With an integrated suite of cloud-based solutions, Proofpoint helps companies around the world stop targeted threats, safeguard their data and make their users more resilient against cyber attacks. Leading organisations of all sizes, including more than half of the Fortune 1000, rely on Proofpoint for people-centric security and compliance solutions that mitigate their most critical risks across email, the cloud, social media and the web. More information is available at **www.proofpoint.com**.