mimecast®

The Connected Human Risk Management Platform

EXPOSING. human risk

Introduction

In our current cybersecurity environment, where threat actors carry snazzy monikers like 'Volt Typhoon' and 'Dark Scorpius', it's unfortunate that everyday users often get overlooked or underestimated in cyber risk assessments.

But ask security leaders about what keeps them up at night—where they feel the most exposed—and it's likely they'll mention threats lurking inside their own organizations.

We aim to flip the script on human risk in this report—expose it in order to reduce our exposure to it. We'll shine the light of data from Mimecast's expansive telemetry on what risky behavior looks like, how often it occurs, and who's engaging in it.

Here's a sample of what we uncovered.

KEY FINDINGS.

48%

Almost half of employees engaged in behaviors that exposed their organizations to cyber risk.

1/3

WEB BROWSING

Of users violated web browsing policies meant to keep them safe.

5%

PHISHING ATTACKS

Expect about 5% of your workforce to fall for phishing attacks each year.

13%

PHISHING EMAILS

Click rates on phishing emails among users averaged 13%. Training reduces that by 25%.

About **1 IN 7** malware-prone employees triggered **10+** events each.

WHO ARE THESE RISKY EMPLOYEES?

Executives, sales, and the board of directors top our list for risk exposure. Read on to discover other “phishy” profiles based on role and tenure.

Benchmarking Risky Behavior

This section measures how often employees engage in behaviors that put their organizations at risk to various cybersecurity threats.

RISK

We focus on three types of risky behaviors: clicking on phishing emails, downloading or executing malware, and violating web browsing policies. These aren't mutually exclusive, of course, and we'll examine rates of recidivism in a later section.

Overall, almost half (48%) of all users engaged in at least one of these behaviors during the timeline of our analysis. Browsing violations occurred most often (36% of users) and malware events were the least common at ~2% of users.

Figure 1: Percentage of users engaging in risky behaviors

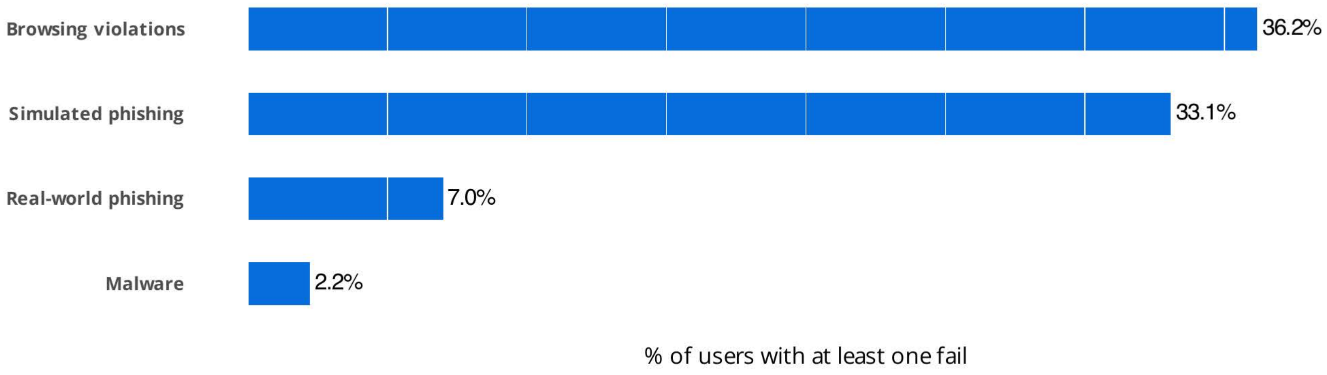

We've included two categories for phishing in the chart, one for clicking on real, malicious phishing attempts and another for simulated phishing exercises run by their organizations to help inoculate them against the real thing.

If you're wondering how well that works, hold that thought—we'll get there. For now, just note that users are less likely to fall for real phish than the fake ones.

Real-world phishing

While phishing isn't the most common according to Figure 1 above, we'll start here because it's arguably top of mind when people think of human risk. And there's good reason for that. The long-running and widely regarded Data Breach Investigations Report from Verizon consistently lists phishing as a top threat action. Cyentia Institute's Information Risk Insights Study found that phishing was among the top three initial access techniques for 18 of 20 sectors.

How often do phishing attacks cross Mimecast? How likely are users to click on them? How many phishing failures should your organization expect in a given year? We reel in answers to those questions and more.

Observed phishing attempts

We've already shown that 7% of all users were hooked by at least one phishing email. But let's back up and establish some prerequisite measures. Over one-third (36.7%) of users never received a real-world phishing attempt during the span of time in which historical event data is available (which differs for each organization and user).

Among users who did receive phishing attempts, the typical rate was approximately six per year, though there is variation in that rate across the user population. This can be seen in figure 2 below (each dot represents 1% of users). About 13% of users received fewer than one phish per year, but 4% of them were targeted with more than 100. Who are those users most targeted by phishing? Good question—we'll tackle that later.

Figure 2: Distribution of phishing attempts per user per year. Each dot represents 1% of users

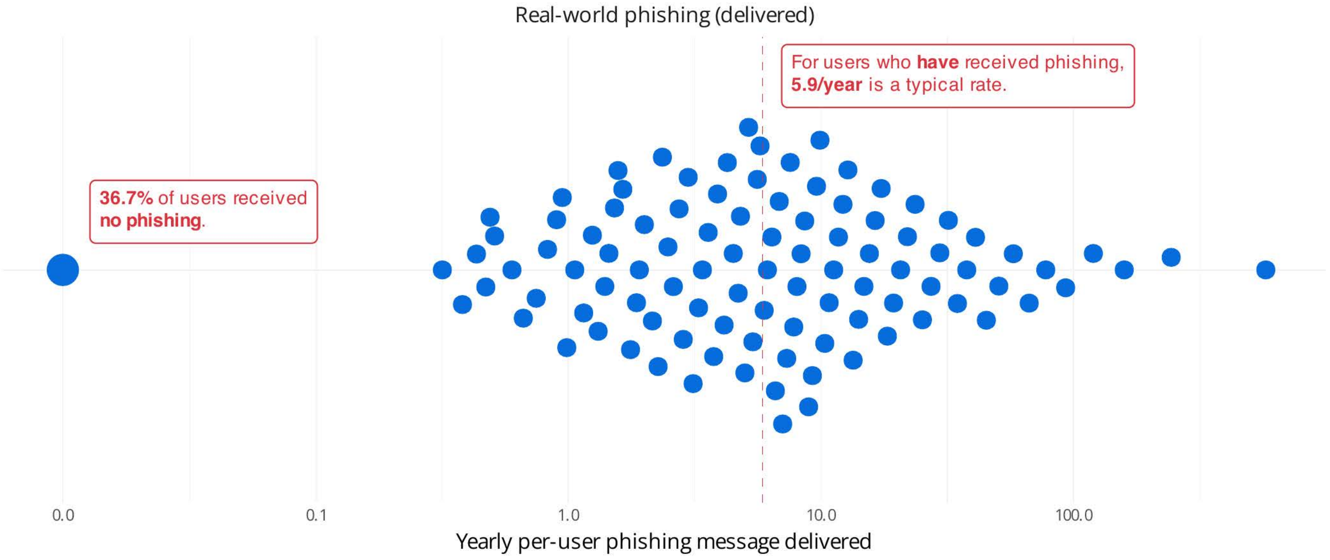

What kind of phish are swimming around?

Mimecast's detections are vast, spanning more than 42,000 organizations around the globe, and as you might imagine, they scoop up all manner of phish. These are examined and categorized into the phishing subtypes you see listed on the left side of Figure 3. We then compare normalized detection rates (per user per year) for each subtype across industries. The shading is relative to columns in the table to help highlight which types of phish are most common to each industry.

Your eye is likely drawn to the dark crimson band for credential harvesting that runs unbroken across all sectors. If you needed more evidence that attackers covet legitimate user credentials as a means of attaining and elevating access into target environments, here it is.

Impersonation is also a common type of phishing across all industries (especially Healthcare and Education) and further corroborates that point. This adds more evidence to the fact that insiders are vectors of attacks far more often than they're the villains behind them.

Note that we've included subtypes for blocked and phishing URLs in a separate tier of the table. We did that because those detections are explicitly triggered by users clicking on phishing messages, resulting in attempts to connect to malicious sites. Thus, they represent an outbound, rather than inbound, view of phishing activity. We chose not to add color shading to further emphasize this difference, but it's worth noting the relatively high rate of detections for blocked URLs across all sectors.

Figure 3: Comparative rates of phishing subtypes detected by sector

		PS	IT	Retail	Sci/Tech	Education	Finance	Manufacturing	Government	Construction	Healthcare
	Abused Fee Fraud	0.008	0.005	0.008	0.006	0.002	0.009	0.005	0.014	0.009	0.003
	Abused Fee Scam	0.057	0.058	0.062	0.061	0.011	0.073	0.038	0.090	0.067	0.024
	Abused Legitimate Services	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	BEC Whaling	0.003	0.004	0.005	0.004	0.000	0.004	0.004	0.000	0.005	0.001
	Credential Harvesting	0.986	1.182	1.074	1.125	0.045	1.491	0.941	0.332	1.683	0.347
	Dating	0.001	0.001	0.002	0.001	0.000	0.001	0.001	0.001	0.002	0.001
	Exploit	0.004	0.005	0.004	0.004	0.000	0.011	0.004	0.001	0.008	0.002
	Fraud	0.099	0.113	0.097	0.077	0.007	0.122	0.087	0.028	0.107	0.042
	Impersonation	0.491	0.893	0.563	0.599	0.041	0.712	0.542	0.134	0.720	0.379
	Low Reputation	0.015	0.012	0.014	0.011	0.000	0.020	0.012	0.002	0.018	0.006
	Malicious File	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	Malspam	0.001	0.001	0.000	0.000	0.000	0.001	0.001	0.000	0.000	0.000
	Monitored Actor	0.202	0.204	0.197	0.201	0.008	0.274	0.180	0.069	0.272	0.065
	Romance Fraud	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.001	0.000
	Sending MTA Detection	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	Suspected Spam	0.050	0.049	0.049	0.048	0.002	0.260	0.040	0.028	0.048	0.015
	Suspicious Message Content	0.001	0.003	0.001	0.001	0.000	0.001	0.001	0.002	0.001	0.001
	Unsolicited Bulk Mail	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Require user action:	Blocked URL	1.669	2.604	1.292	1.066	0.142	2.054	1.431	0.491	1.568	0.700
	Phishing URL	0.012	0.009	0.009	0.004	0.001	0.075	0.011	0.005	0.016	0.009

Phishing click rates

Receiving phishing emails is one thing; falling for them is another thing entirely. According to our analysis, 89% of users who received real-world phishing never clicked on any of them. Bravo!

Of those users who did take the bait, the typical likelihood of clicking was 12.5%, though once again, we see a wide disparity among them. We observed users with click rates as low as 0.1% and others who fell hook, line, and sinker for every phishing attempt cast their way.

Figure 4: Distribution of click rates among users for phishing attempts

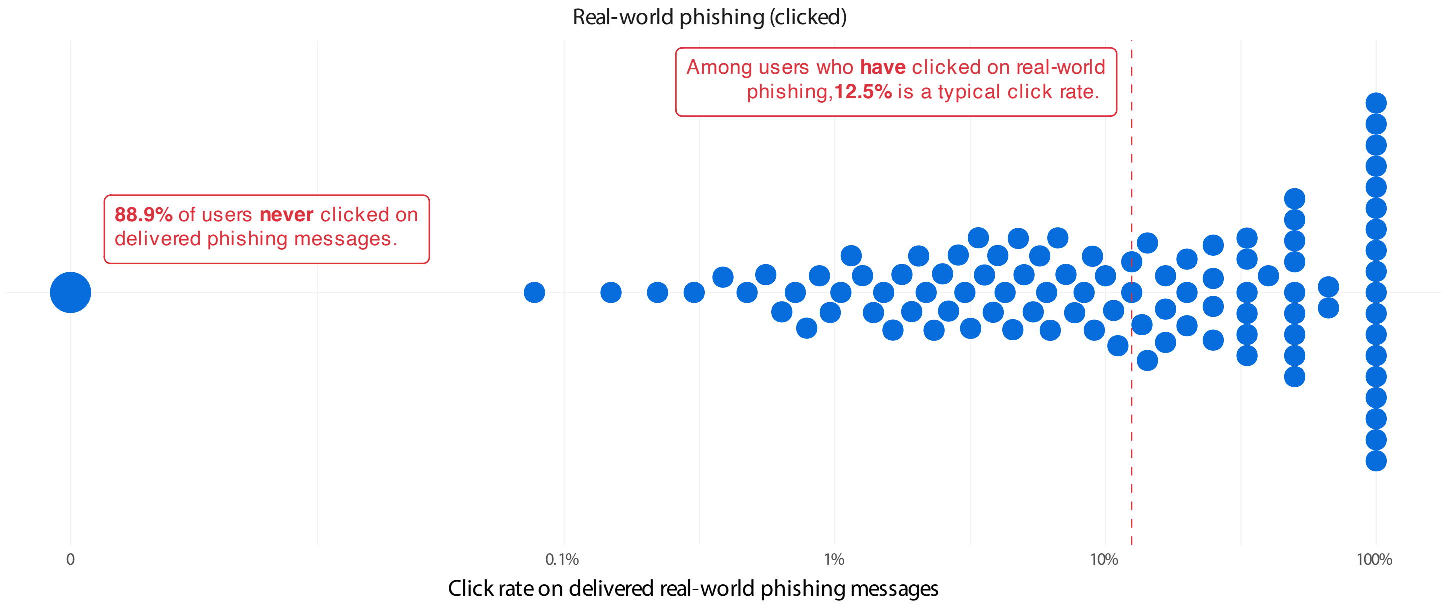

Expected frequency of successful phish

How many phishing emails will net clicks in your organization in the next year? Well, that's tough to answer without knowing more about your particular organization. But what we can do is apply some math (specifically, Empirical Bayes) to our data on historical delivery and click rates to model the expected frequency of successful phishing attacks.

Using that model, we can make some projections for a 1,000-person organization. Just under 50 users (48) will click on at least one phishing message per year. Nine employees will fall for two or more, and one poor user will be hooked more times than that. Figure 5 visualizes this information.

Figure 5: Modeled frequency of successful phishing attacks per anum

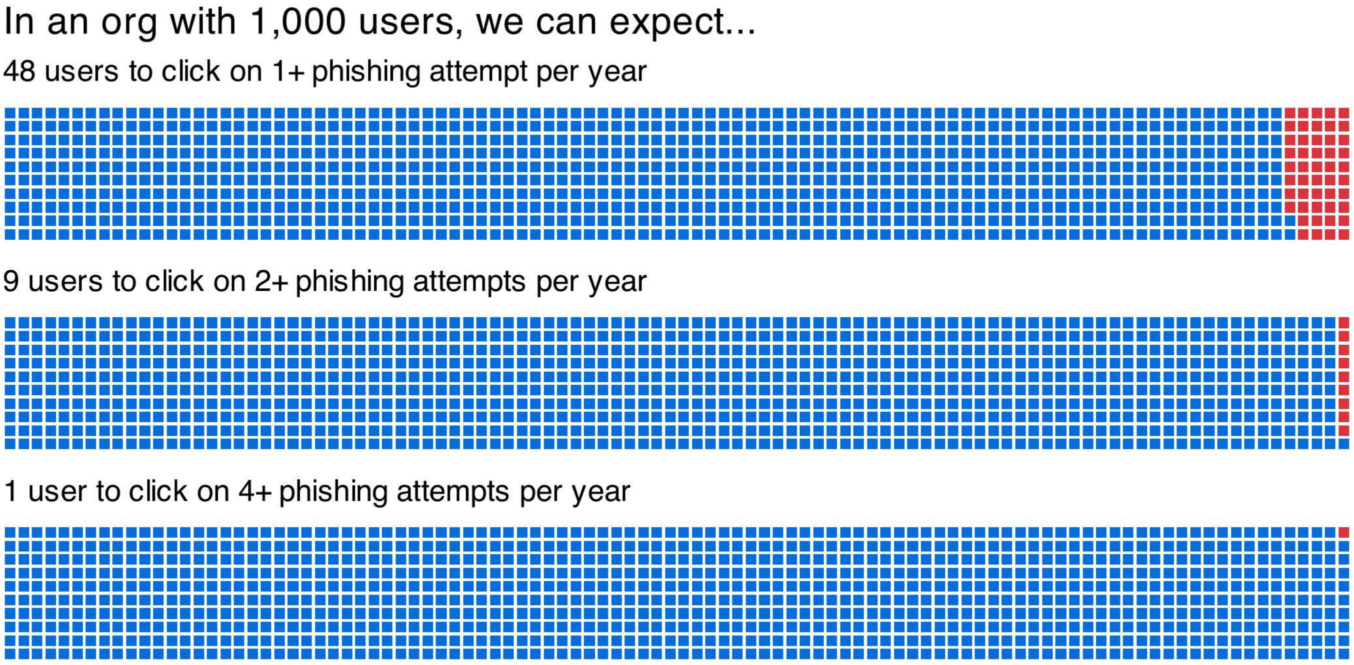

Can training help kick the click?

There is some evidence that employees can be trained to kick their clicking habit to a certain degree, but the evidence also warns that “clickers gonna click.”

We examined phishing click rates among users before and after completing training. We observed very different effects depending on the employee’s propensity to click. Those who already exhibited low click rates showed no additional improvements in the months after a training session. But those with a tendency to click averaged a 25% reduction in click rates.

Organizations should consider augmenting training and intervention for their employees most prone to click. This can entail more timely intervention tied back to real-world clicks and risky events.

These results suggest that, while training definitely won’t “kick the click” entirely out of your organization, it can, at least, help curb that behavior among your riskiest users. They also hint that a more targeted, tailored approach to training and other interventions will likely meet with greater success than following the same script for everyone.

Figure 11: Comparison of average reduction in phishing click rates after training

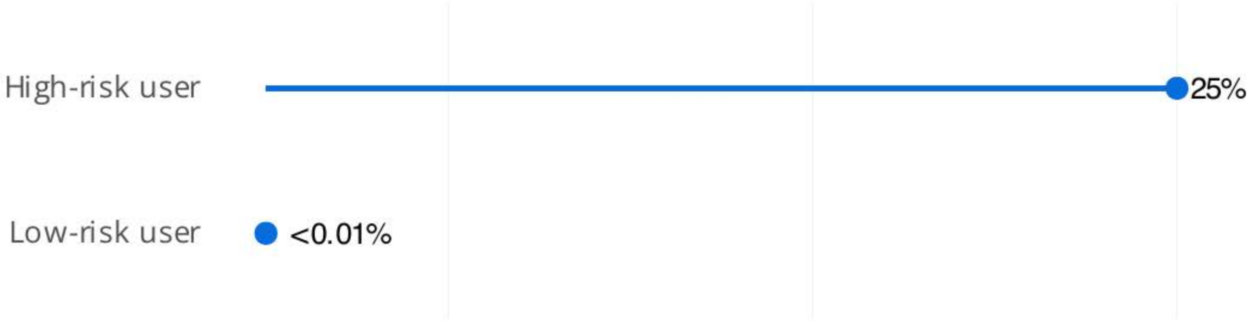

Malware events

Malicious software, or malware, is the multi-tool of the cybercriminal world. It offers the ability to communicate remotely, issue commands, gain backdoor access, find and exfiltrate data, destroy systems, erase evidence, and much more.

While attackers increasingly try to “live off the land,” using existing tools for their illicit activities, getting employees to download and/or execute malware is still a very common tactic. Thus, it undoubtedly constitutes risky behavior that organizations want to avoid.

Let’s see how they’re doing.

Observed malware encounters

Good news is often hard to find in cyber threat reports, so let's start by recognizing that nearly all (98%) employees made it through our sample time period with a spotless record for malware events. That speaks to the many anti-malware defenses that exist between users in modern organizations and the malware-ridden internet around them.

But the 2% of users who did download or execute malware obviously can't be ignored. The typical rate of occurrence among them was about 1.5 malware events per year. As we saw with phishing, there's a lot of variation in that rate in Figure 6. About one in seven employees were solely responsible for triggering 10 or more malware events.

Figure 6: Distribution of malware downloads/executions per user per year

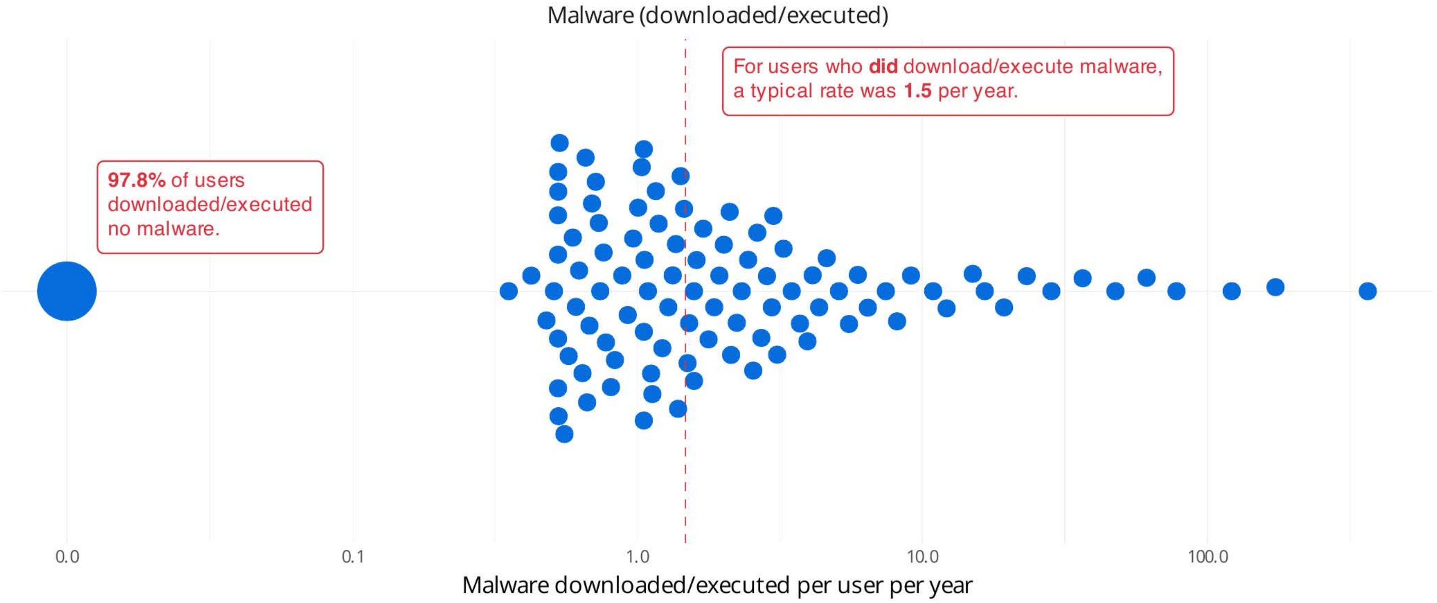

What kind of malware is milling about?

Mimecast sensor collections offer some additional granularity on the types of malware employees are encountering as they carry out their activities. The top category for most sectors is malware samples associated with known threat actors. Those interested in examples of specific malware used by various threat groups will find them aplenty in the MITRE ATT&CK site.

The oddball in the Monitored Actor dominance is the Finance sector. Exploits of vulnerable software and hardware top the list for that sector, possibly because financial services firms tend to have more mature controls in place. Vulnerabilities open holes in those otherwise strong defenses that can be quickly weaponized and exploited by malware.

In general, there's far more variation among organizations than industries. While it is true that the majority of manufacturing firms encounter malware at a higher rate than educational institutions, the overlapping distributions serve as a reminder that's not always the case.

Figure 7: Comparative rates of malware subtypes detected by sector

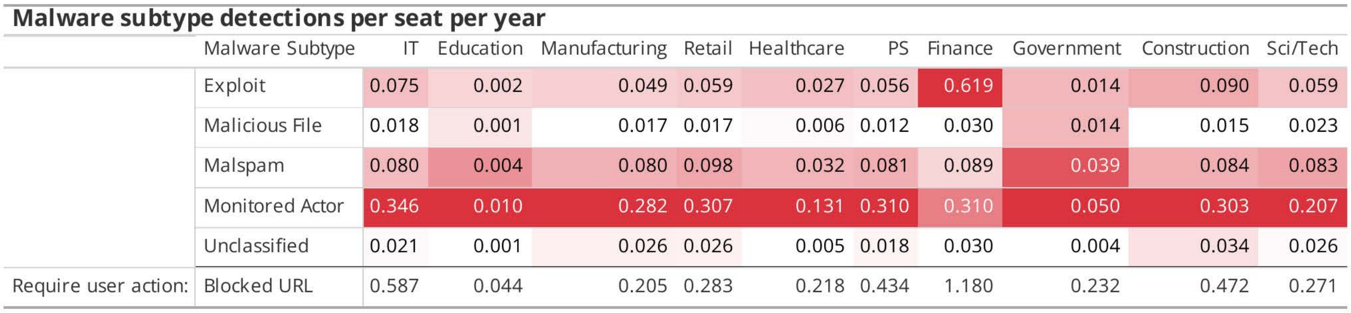

Malware subtype detections per seat per year

	Malware Subtype	IT	Education	Manufacturing	Retail	Healthcare	PS	Finance	Government	Construction	Sci/Tech
	Exploit	0.075	0.002	0.049	0.059	0.027	0.056	0.619	0.014	0.090	0.059
	Malicious File	0.018	0.001	0.017	0.017	0.006	0.012	0.030	0.014	0.015	0.023
	Malspam	0.080	0.004	0.080	0.098	0.032	0.081	0.089	0.039	0.084	0.083
	Monitored Actor	0.346	0.010	0.282	0.307	0.131	0.310	0.310	0.050	0.303	0.207
	Unclassified	0.021	0.001	0.026	0.026	0.005	0.018	0.030	0.004	0.034	0.026
Require user action:	Blocked URL	0.587	0.044	0.205	0.283	0.218	0.434	1.180	0.232	0.472	0.271

Expected frequency

We used the same basic approach described for phishing to model the frequency of malware events to derive a normalized estimate. In a 1,000-person organization, we expect 14 to download or execute malware. Seven of those employees will trigger malware events on a monthly basis, and four will find their way into weekly encounters with malicious software.

If that seems like a small number of users behind a large number of events, you've caught onto an important aspect of human risk: it's not evenly distributed across all employees. We'll pull more on that thread in the next section, but let's first finish up our trio of risky behaviors with browsing violations.

Figure 8: Modeled frequency of malware downloads/executions

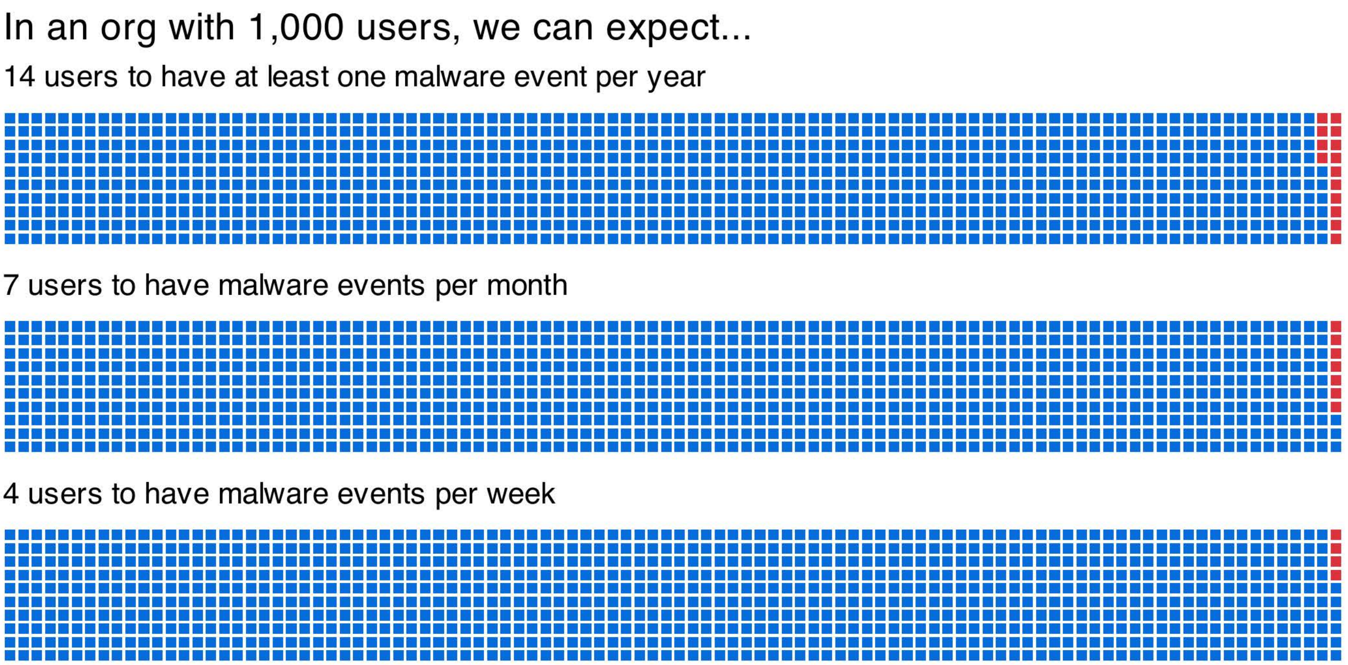

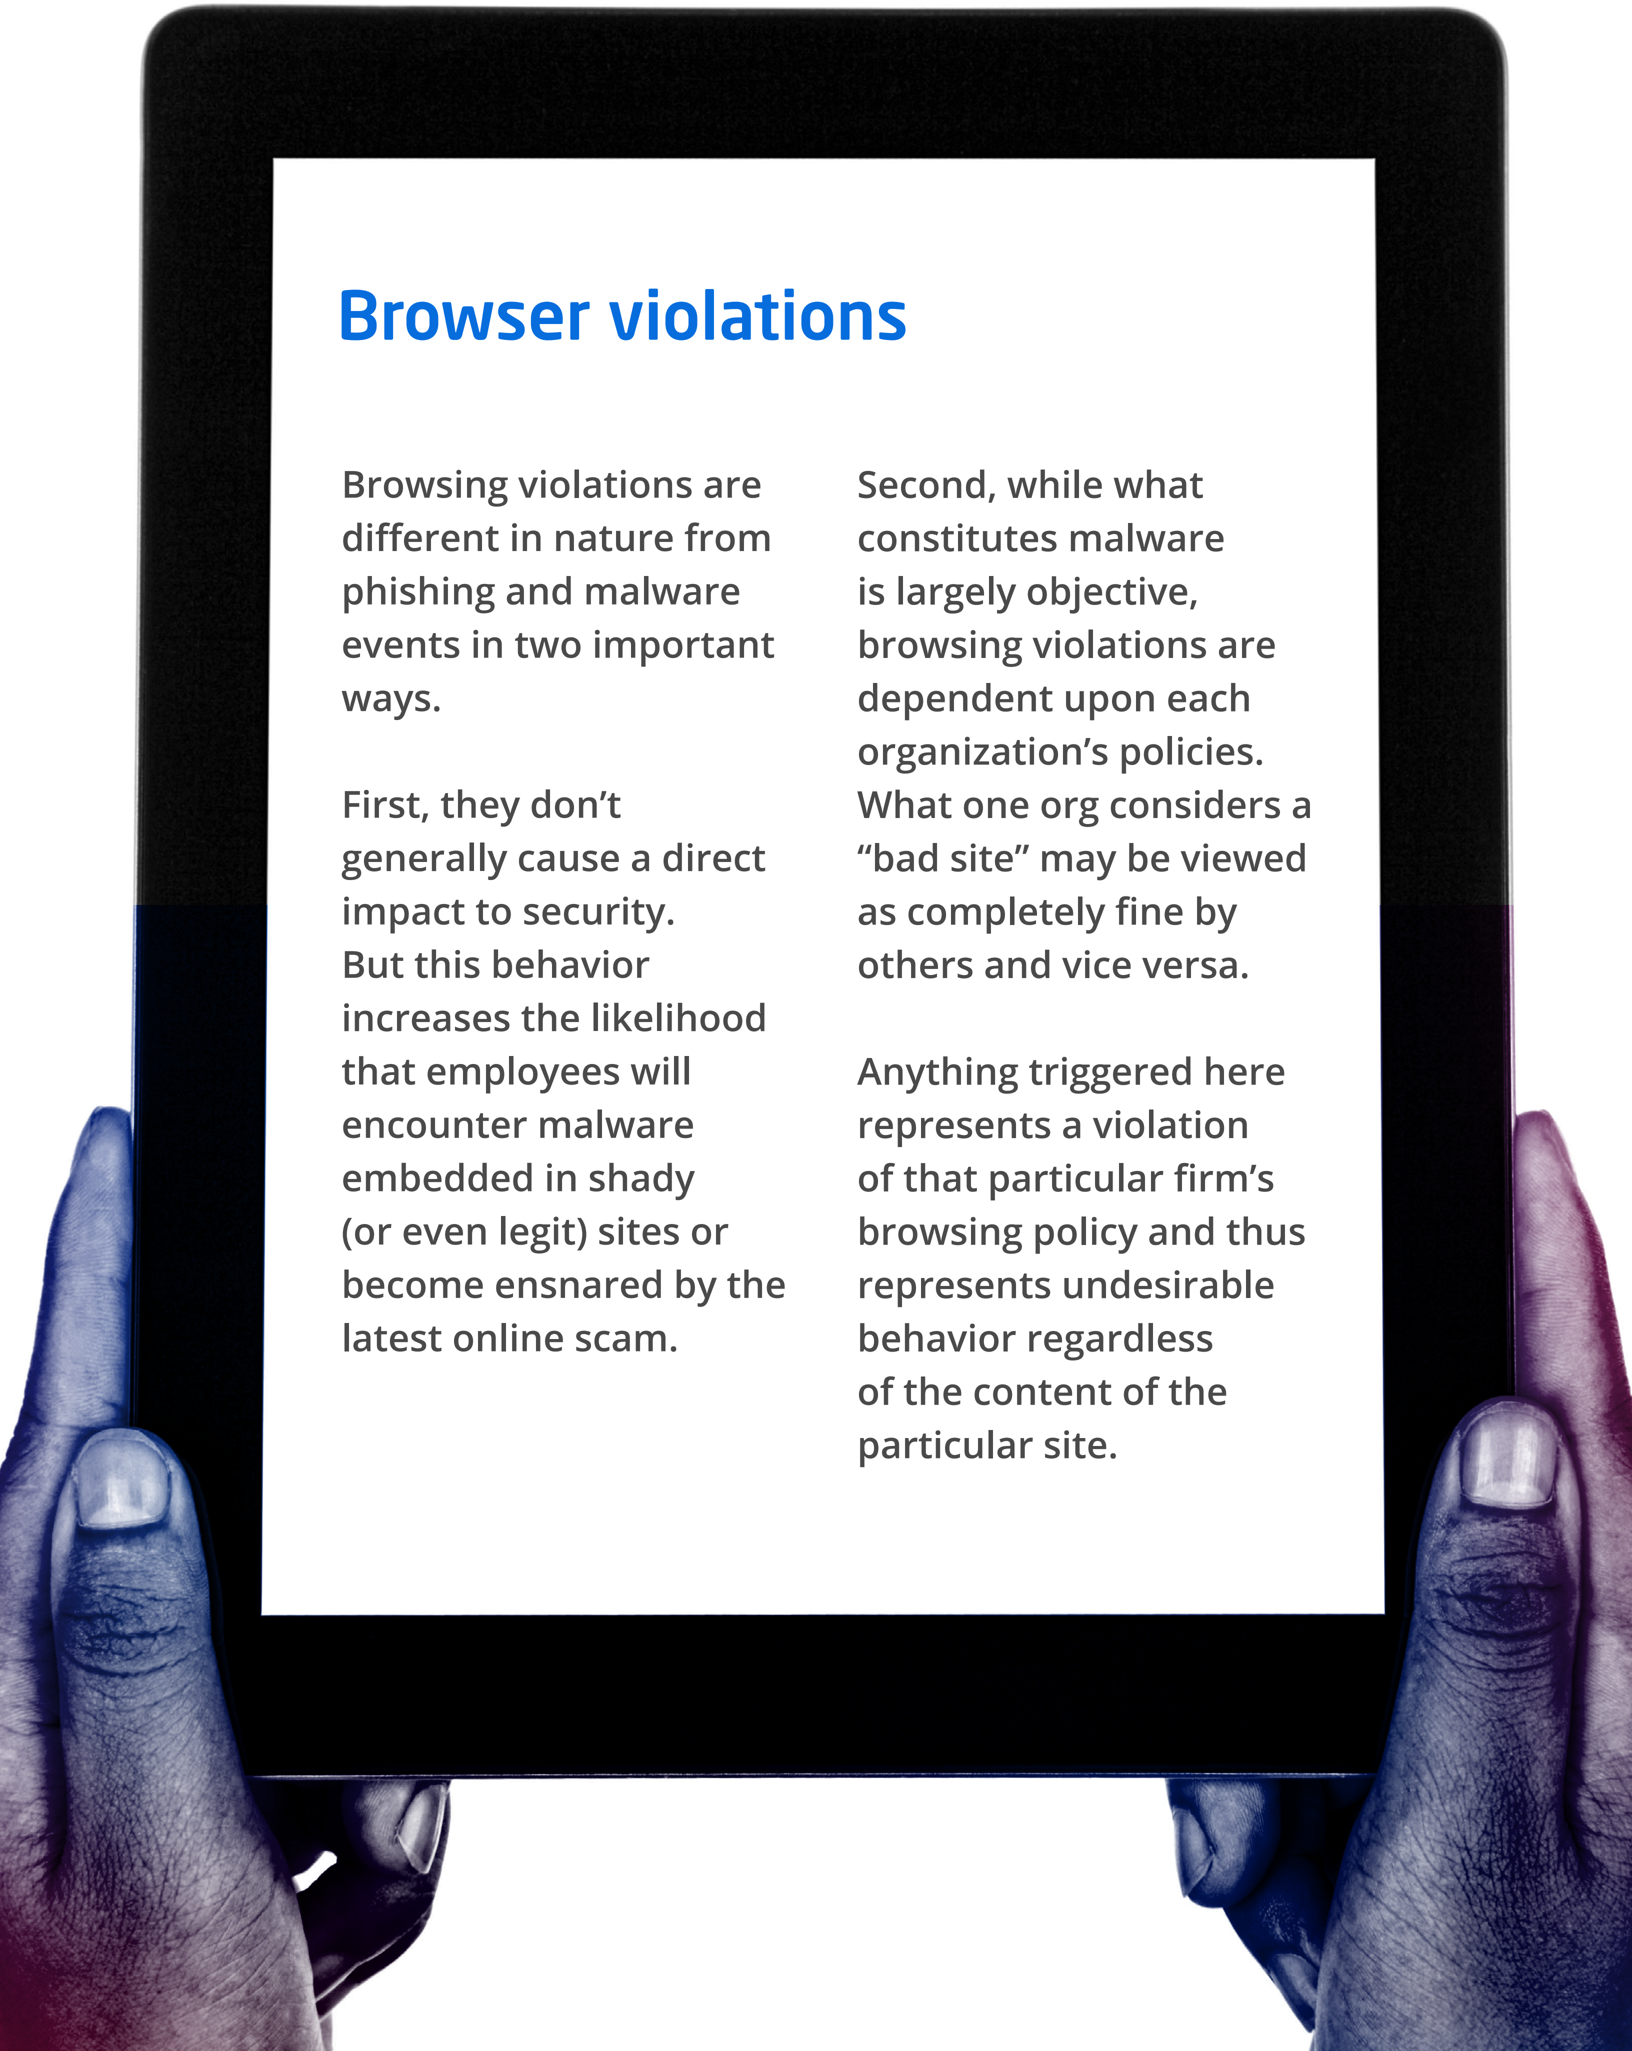

Browser violations

Browsing violations are different in nature from phishing and malware events in two important ways.

First, they don’t generally cause a direct impact to security. But this behavior increases the likelihood that employees will encounter malware embedded in shady (or even legit) sites or become ensnared by the latest online scam.

Second, while what constitutes malware is largely objective, browsing violations are dependent upon each organization’s policies. What one org considers a “bad site” may be viewed as completely fine by others and vice versa.

Anything triggered here represents a violation of that particular firm’s browsing policy and thus represents undesirable behavior regardless of the content of the particular site.

Observed frequency

As seen back in Figure 1, browsing violations are a lot more common than phishing and malware events. Users who engage in this behavior are still in the minority, however—64% of them never triggered violations in our time period of observation. Employees who did log browsing violations averaged under two per year (see Figure 9 for full distribution).

Figure 9: Distribution of browsing policy violations per user per year

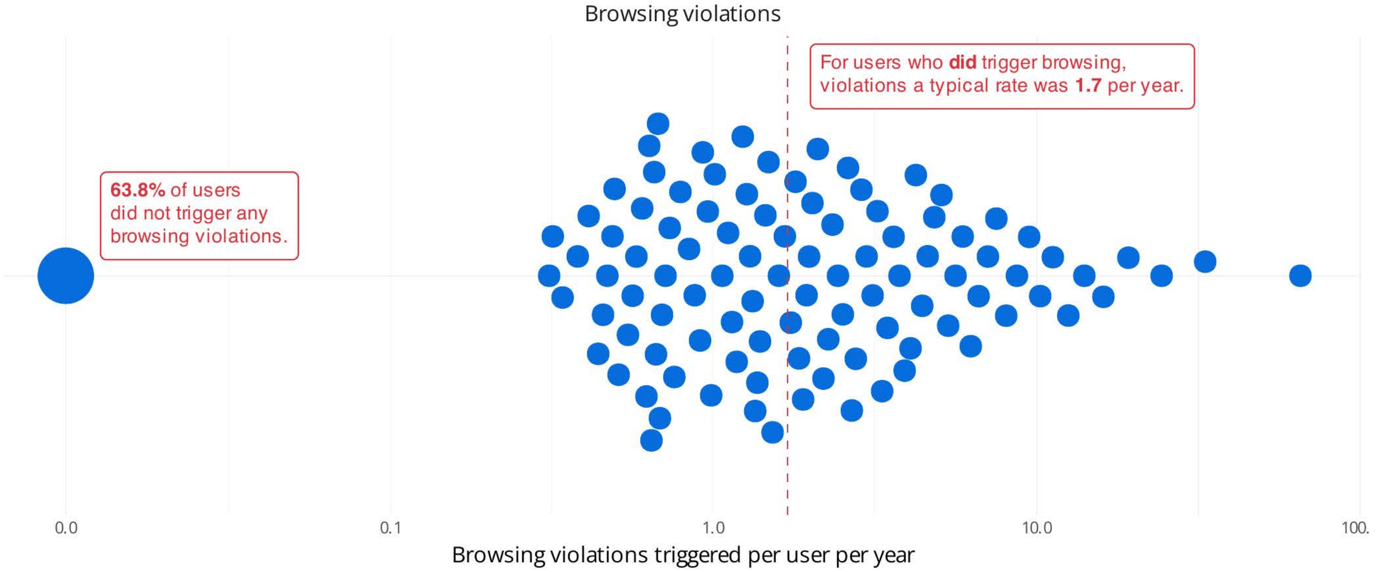

Expected frequency

Since there’s no intermediate step like clicking on a phishing link or executing a malware attachment to measure for this behavior, we’ll jump straight from the observed to the expected frequency. We’ve applied the same approach from the previous two behaviors to model the frequency of browsing violations to derive a normalized estimate.

In an organization of 1,000 employees, we could expect 244 users to violate web browsing policies in a given year. Sixteen of those employees are likely to generate browsing violations on a monthly basis.

Figure 10: Modeled frequency of browsing policy violations

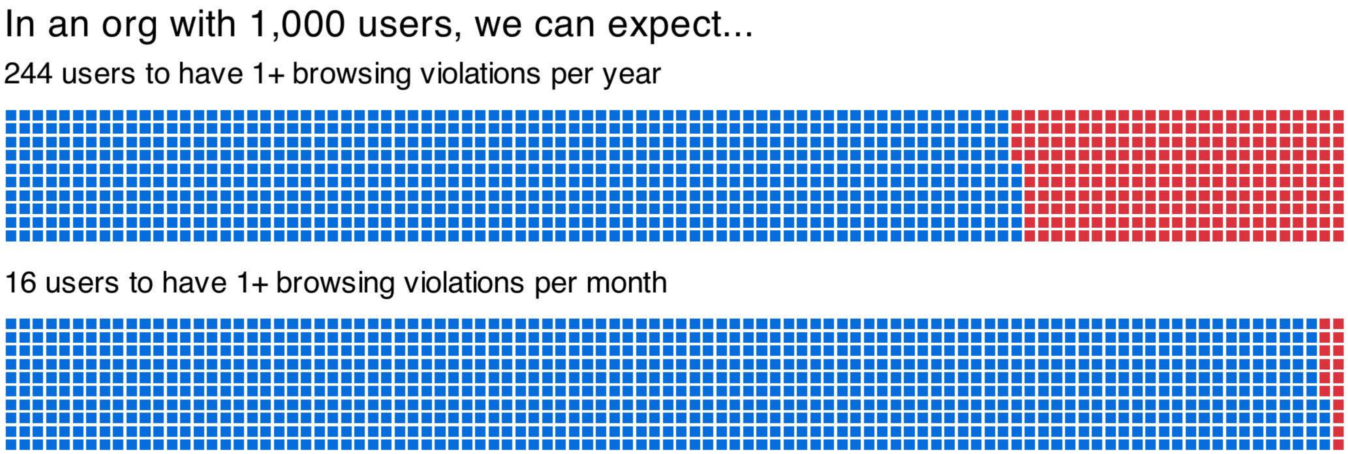

When Risk Becomes Habit

Reading through the last section, you may have noticed a common and important trend: the majority of employees refrain from risky behaviors but a subset make them a habit. Some repeat offenders do the same thing over and over again (e.g., often lured by phishing), while others engage in multiple undesirable behaviors (get phished, download malware, etc). Let's take a closer look at the various forms of high-risk users.

HABIT

Repetitive risky behavior

The next two charts demonstrate that a small number of users can be responsible for an abnormally large share of risky behavior. That's a point of concern for anyone managing human risk. On the positive side, this presents an opportunity to have a huge impact on risk exposure by changing the behavior of a few individuals.

We chose not to show it here, but browsing violations exhibit a similar, though not as pronounced, pattern of dominance by the few. The upper 5% of promiscuous browsers generated 62% of all browsing policy violations. If we look across all three risky behaviors (phishing, malware, and browsing), 5% of users are behind 75% of all detected events.

75%

Figure 12: Phishing events among users

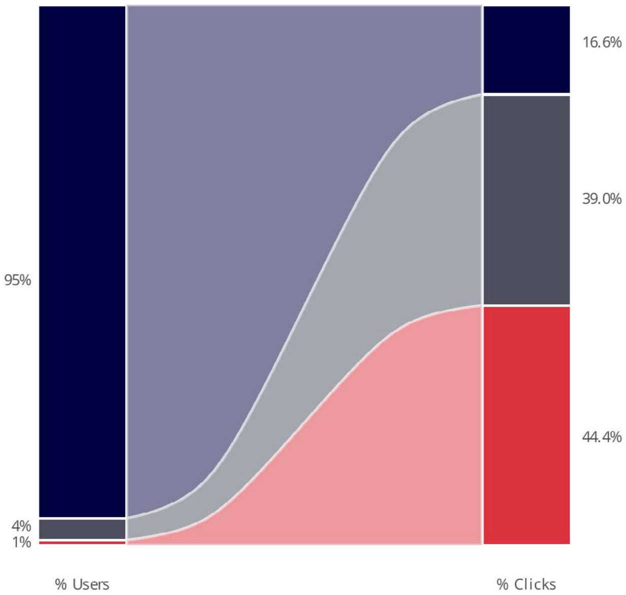

- Just 1% of users are behind 44% of all clicked phishing emails.
- 5% of users are responsible for 83.4% of all clicks.
- The remaining 95% of users collectively account for less than 17% of successful phishing attacks.

Figure 13: Malware events among users

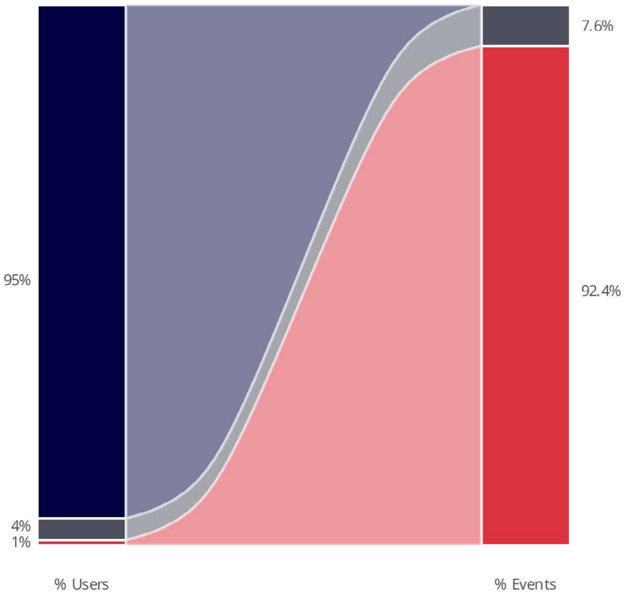

- 1% of users are behind 92% of all malware events!
- 5% of users are responsible for ALL malware events. The remaining 95% had a clean record.
- Malware is far more "lopsided" than the other event types.

Multiple risky behaviors

Having established that a few users tend to cause the bulk of risky events, one may wonder if the same subset of users repeatedly falling for phishing schemes are also downloading malware and violating browsing policies at high rates. Let's take a look.

Among the 48% of employees who engaged in some form of risky behavior, most managed to keep it to just one type (Figure 14). The percentage of users flagged in two behavior categories drops to 13%. Less than 1% transgressed in three or more risky behaviors.

Figure 14: Percentage of users engaging in multiple risky behaviors

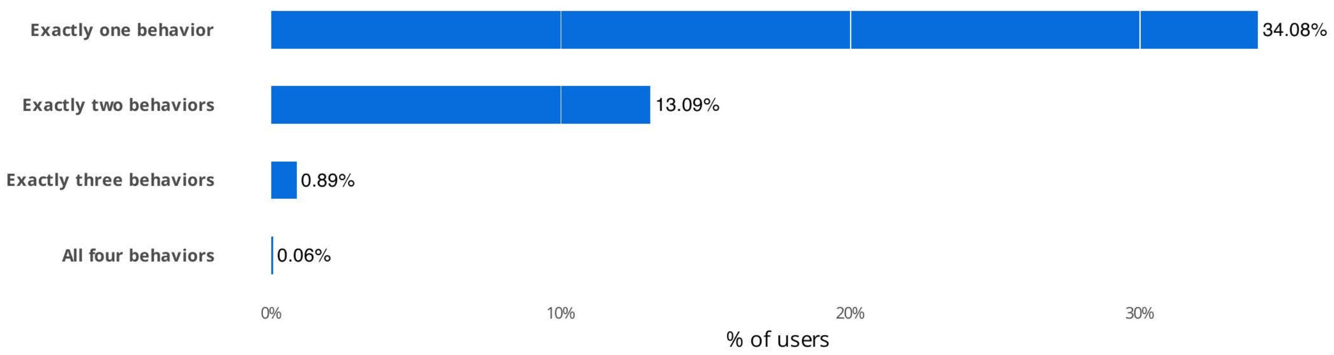

Now, you're perhaps wondering which types of misbehaviors tend to occur in tandem. At least, that's where our minds went next, leading to creation of the "UpSet" diagram on next page.

Figure 15 presents a breakdown of the 48% of all employees in our dataset who engaged in some form of undesirable behavior (the bar for the 52% who had a clean record is omitted). Readers may find certain intersections of behaviors more or less interesting for different reasons, so we'll highlight something that stood out to us and leave you to glean your own takeaways.

It's not surprising that the largest bar that includes real phishing is the combination of users who failed real and simulated phishing (3.39%). What is interesting is that the next largest bar including real phishing is the set of users who failed real phishing but nothing else (1.05%). In fact, all the combinations that involve users who failed real phishing but not simulated phishing amounts to 1.38% of users. Granted, that's not a huge percentage, but it's not ignorable either. Why are these employees slipping through the simulations? Could it be that simulated phishing messages are too tricky?

Figure 15: Overlaps in risky behaviors among users

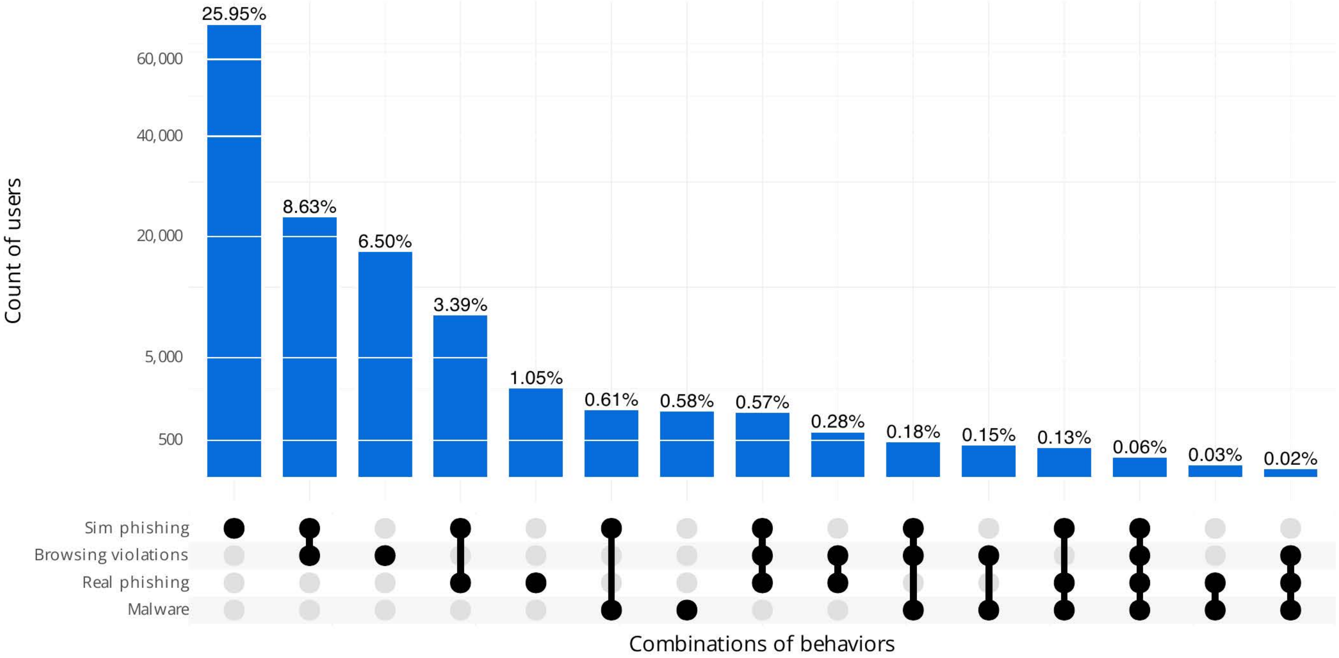

Are simulated phish too tricky?

This question posed at the end of the previous paragraph wasn't just rhetorical. We'll pick it up for a closer look here. Ideally, regular simulated phishing trials would root out all employees prone to taking the bait so they learn not to bite at the real deal. Over time, we'd expect click rates for simulated and real phishing attempts to be similar if tests mimicked the actual attacks. But this is not at all what we see in the data.

Per Figure 16, click rates for simulated phishing trials are much higher than for real-world phishing attacks. So much so, in fact, that their distributions hardly overlap (which statisticians would interpret as indicating these are fundamentally different things).

A possible explanation of what we're seeing here is that real phish are easier for employees to spot than their simulated cousins. At the very least, they don't appear well calibrated. We can't help but wonder if that disparity could be misleading employees about what real phishing messages look like, enabling attackers to slip in through the sims.

Figure 16: Comparison of user click rates between real vs. simulated phishing emails

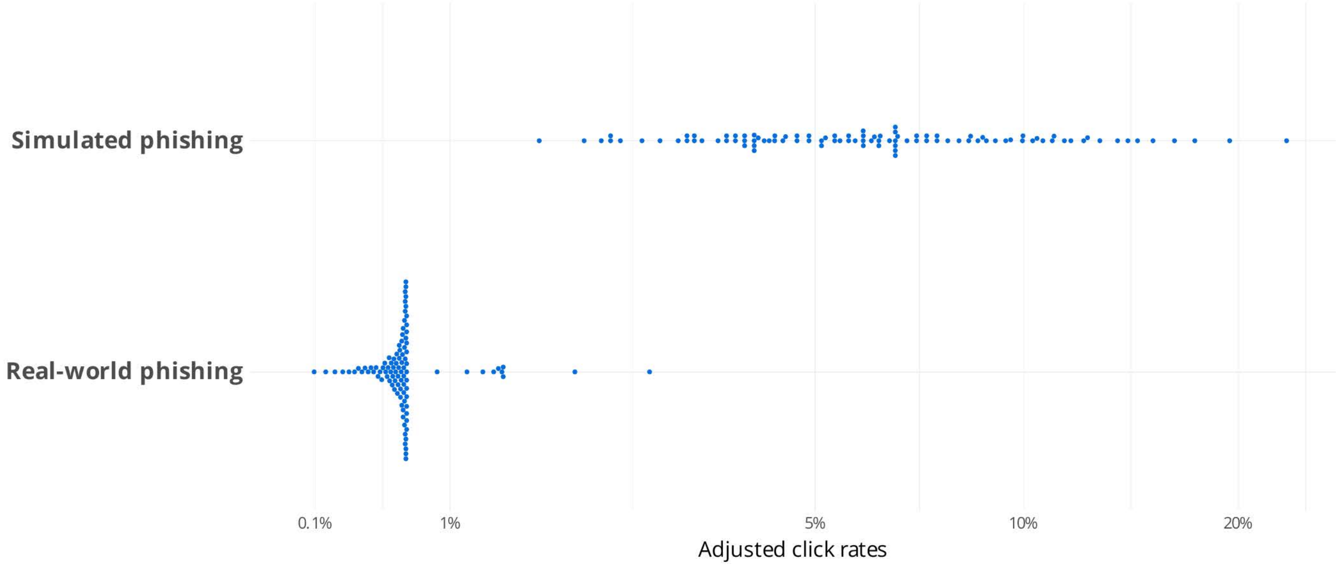

Risky users: targeted or tricked?

The last two subsections have focused on behaviors that cause some employees to represent higher risk than others. But is human risk entirely based on what users do? Or is there also an aspect of who they are that makes one user's risk profile different from another?

Mimecast's phishing telemetry provides a useful lens through which to study this question because we can separate receiving phishing emails (targeting) from the act of being tricked into clicking on them. We'll start with a role-based comparison of these measures.

According to Figure 17, managers are targeted by phishing attacks far more often than regular employees or contractors. That probably reflects a more public persona and higher levels of access/influence. That said, managers are the least likely to click on those phishing messages. Even so, the last column shows they have the highest expected rate of successful phishing incidents (per user, per year). It's important to note, though, that the rate of apparent targeting is what elevates managers' risk profile. That suggests shielding them from those attacks could be more effective than mandating additional training.

Figure 17: Comparison of phishing risk metrics among organizational roles

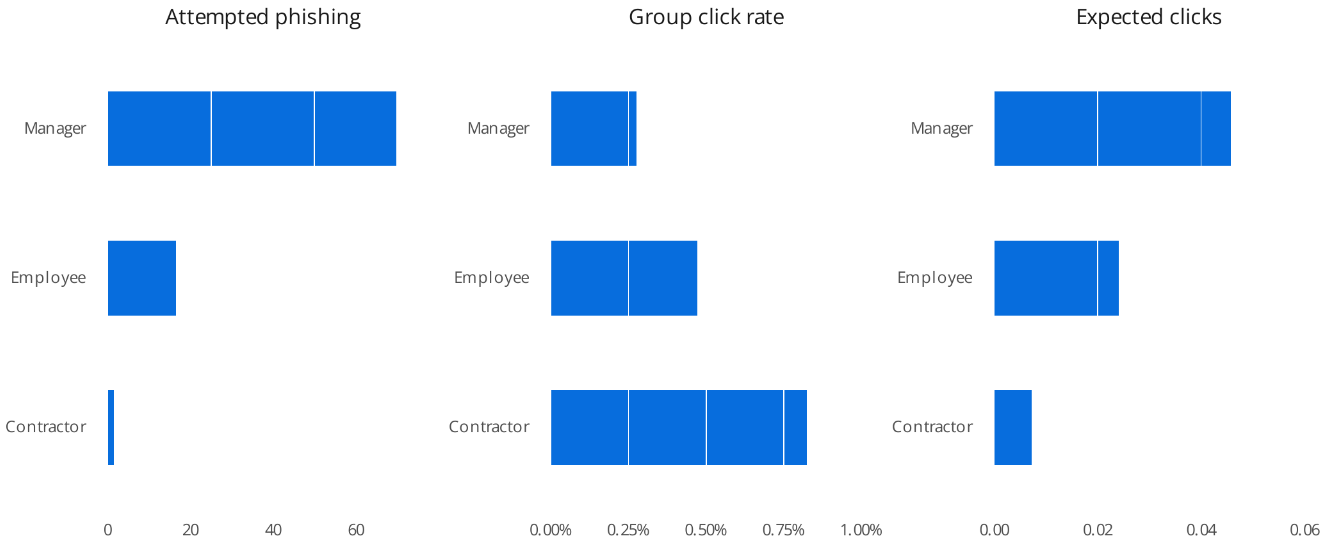

Figure 18 takes a more detailed look at risky roles by comparing different organizational departments or functions. Based on the prior chart, it's not surprising to see executives receive the most phishing emails. But sales and the board of directors are right up there with them. All of these tend to be very public-facing roles, which lands them on the phishers' radar. Even though these roles tend to have low to average click rates, their probability of being successfully phished exceeds all others.

Lab employees serve as a great example of the "targeted vs. tricked" distinction. They receive the fewest phishing emails but are the most likely to click on them. Customers exhibit a similar pattern. This makes them ideal candidates for some well-designed training or phishing simulations to lower those click rates and reduce their overall risk profile.

Figure 18: Comparison of phishing risk metrics among departments

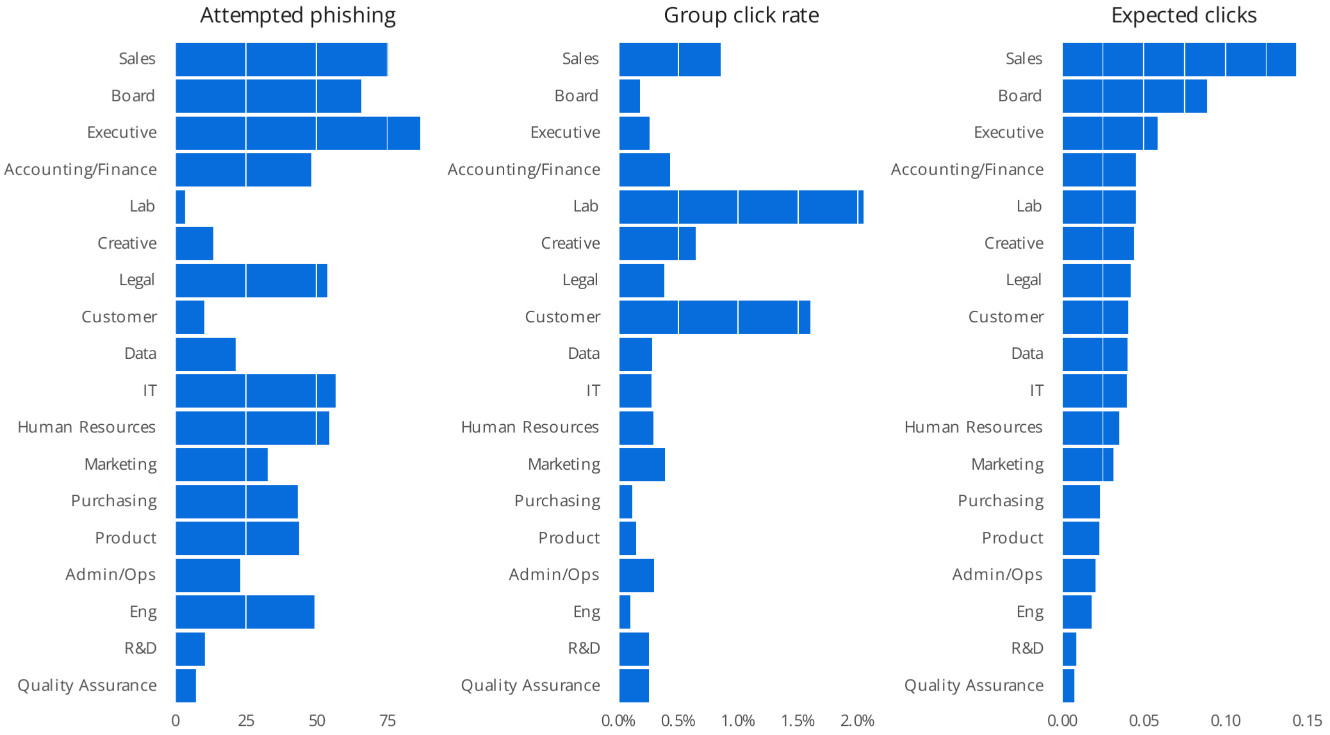

Last but not least, let's see how tenure shapes an employee's risk profile. The short story here is that the longer you're around, the more you're phished. That probably has a lot to do with corporate email addresses being added to more and more cybercriminal contact lists over time.

Click rates show the opposite trend; the newest employees are most readily duped. In terms of successful attacks, two years of tenure appears to be a breakpoint where the expected risk doubles. As with managers, this is due to elevated targeting of more tenured employees, and yet more training is unlikely to offset the risk.

Figure 19: Comparison of phishing risk metrics by tenure with current employer

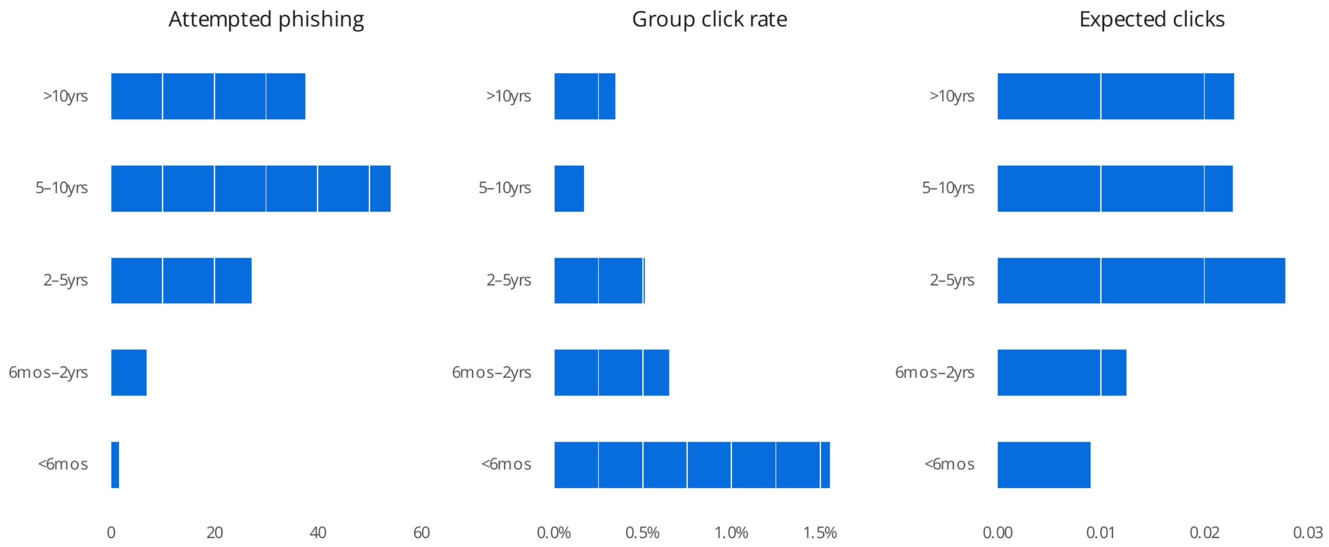

CON CLU SION

Despite the “cyber” prefix, cybersecurity starts and ends with people. Human behavior remains a significant vulnerability in even the most secure environments. The data underscores that nearly half of employees exhibit risky behaviors that expose their organizations to phishing, malware, and other cyber threats. This persistent human risk poses a challenge to cybersecurity leaders, but it also presents an opportunity.

Cybersecurity leaders must therefore adopt a proactive, human-centric approach to managing risk. This requires moving beyond basic awareness training and focusing on behavioral change through targeted, continuous education and reinforcement.

As shown in the Mimecast study, repetitive risky behaviors are often concentrated within a small percentage of employees. This small group accounts for the majority of security incidents. **Tailored interventions for these high-risk users are critical.**

Cybersecurity is no longer just about preventing external breaches but managing the risks that originate from within. By understanding and mitigating human risk, cybersecurity leaders can build stronger defenses and reduce the chances of costly security incidents in the future.

Cybersecurity leaders should implement a data-driven approach to identify, engage, and educate these individuals.

This includes:

1

Leverage Risk-Specific Training and Intervention

Use advanced behavioral analytics to deliver targeted training to employees exhibiting repeated risky behaviors, especially those in roles more susceptible to phishing attacks, such as executives and sales teams.

2

Enhance Risk Visibility

Ensure that user-based risk analysis accounts for more than just phishing simulation exercises. Discrepancies in phishing simulations can diminish the effectiveness of these exercises. Security teams should consider including other behavior-based data when assessing human risk.

3

Develop Role-Based Protections

Since certain roles (e.g., executives, sales, board members) are more heavily targeted, deploy additional layers of protection and monitoring for these individuals. This includes reducing their exposure in public-facing situations.

4

Adopt a Holistic Human Risk Management Framework

Integrate security technologies with a human-centric strategy that fosters continuous engagement and accountability. Mimecast's AI-powered, API-enabled Human Risk Management platform is a perfect example of how technology can be used to elevate visibility, offer strategic insights, and take decisive action to reduce risk.

About Mimecast

Mimecast is an AI-powered, API-enabled connected Human Risk Management platform, purpose-built to protect organizations from the spectrum of cyber threats. Integrating cutting-edge technology with human-centric pathways, our platform enhances visibility and provides strategic insight that enables decisive action and empowers businesses to protect their collaborative environments, safeguard their critical data and actively engage employees in reducing risk and enhancing productivity. More than 42,000 businesses worldwide trust Mimecast to help them keep ahead of the ever-evolving threat landscape. From insider risk to external threats, with Mimecast customers get more. More visibility. More insight. More agility. More security.

www.mimecast.com

About Cyentia Institute

The Cyentia Institute is a research and data science firm working to advance cybersecurity knowledge and practice. Cyentia pursues this goal through data-driven studies like this one and through a growing portfolio of analytic services.

www.cyentia.com