Beyond the Click:
Experts Weigh in on the Phishing Training Debate

Authors:

Shanée Dawkins, National Institute of Standards and Technology

Julie Haney, National Institute of Standards and Technology

Abstract

The efficacy of phishing simulation training is a hotly debated topic. Three cybersecurity experts from academia and industry discuss the reasons why many phishing training programs fail to change behavior, the importance of meaningful measures of effectiveness, and strategies for improving organizational resilience against current and future phishing threats.

Introduction

Phishing awareness training has evolved beyond traditional compliance-based training (e.g., information dissemination) to accommodate different modes of learning. A common, interactive method used by organizations today is embedded phishing awareness training, in which simulated phishing emails are periodically sent to employees to train them on what to look for in real-world phishing attacks. Those who fall for the phish by clicking on a link in the email or downloading an attachment are notified that the email is a phish and may receive additional training, for example, tips on how to spot phishing emails in general, or guidance on awareness of that particular type of phish. Conversely, those who report the phish to their organization may be rewarded for their vigilance, for example, receiving a virtual badge or ribbon on their internal account reserved for top reporters.

While these embedded programs have been widely adopted, their effectiveness in achieving behavior change has been debated. Some training program coordinators claim significant reduction in click rates and increased reporting of suspected phishing emails as a result of embedded training [1]. However, research findings from large-scale experiments in organizational settings suggest that these programs may not meaningfully improve real-world human detection of phishing emails [3][4].

To further the conversation around this debate, Shanée Dawkins and Julie Haney from the National Institute of Standards and Technology (NIST) Human-Centered Cybersecurity program facilitated a virtual roundtable with three experts in phishing awareness training: Dr. Grant Ho, an academic researcher; Dr. Chris Fennell, an industry researcher; and Mr. Cary Johnson, an industry practitioner. The roundtable conversation was transcribed, condensed, and edited by the authors with the approval of the roundtable experts.

Published in IEEE Security & Privacy Magazine July/August 2026 issue

Roundtable Experts

Chris Fennell is a distinguished engineer and technical lead on the Research Science and Hardware Intelligence team in InfoSec at Walmart. He received his PhD from Michigan State University, and his research focuses on a variety of security projects from human-computer interaction to systems and networking. His work has led to contributions across several publications, patents, and policy efforts. He is also actively involved in both academic and practitioner research communities and currently co-chairs the Workshop on Security Operations Center Operations and Construction (WOSOC) housed at the Network and Distributed Systems Security (NDSS) Symposium.

Grant Ho is an assistant professor in computer science at the University of Chicago. His research focuses on building a data-driven foundation for protecting organizations and users. Previously Grant was a postdoctoral fellow at University of California, San Diego and received his PhD from University of California, Berkeley. He has published multiple papers on phishing, email security, and authentication, which have been recognized by the 2017 Internet Defense Prize and three distinguished papers awards across top cybersecurity conferences, including USENIX Security Symposium and IEEE Symposium on Security and Privacy.

Cary Johnson is the co-founder of Phishbusters and an independent expert in phishing training effectiveness. He developed an innovative baseline-driven methodology that operates as an independent measurement layer alongside existing awareness and human risk management platforms, enabling organizations to track changes in their phishing resilience over time. By focusing on internal baselines rather than vendor-defined metrics or external benchmarks, he brings a distinct perspective on what is actually working and what is not, for organizations seeking to mature beyond compliance through objective, organization-specific measurement.

Are current phishing awareness and training efforts failing to translate into long-term behavioral change? Why or why not?

Dr. Grant Ho: There is a lot of controversy around phishing awareness training. In our study, what we tried to understand was whether two types of training led to significant increases in anti-phishing behavior [3]. One was annual cybersecurity awareness training that many people have to do at their organizations and the other was embedded phishing awareness training where employees are sent simulated phishing emails once a month or every so often, and people who fail get additional training. The results of this study show that there is some small statistical improvement that people get; people are 2% less likely to fall for phishing if they take some form of training compared to people who get no training at all. So there are some benefits. On the other hand, our study suggests that phishing lures (simulated phishing emails) used in embedded training programs can be very effective. We were able to consistently get up to 10 to 30% of employees to fall for these phishing lures, or at least click on the links. This suggests that phishing training gives you some benefit, but it's really not doing a whole lot in comparison to the efficacy of some phishing lures. The small benefit is probably not enough to make a huge improvement in the face of current and future phishing attacks.

Mr. Cary Johnson: I believe Grant's research was based more on compliance-focused deployments. Compliance-based training often doesn't work well when it comes to addressing phishing because it is not designed to create or measure behavioral change. Meeting compliance requirements sit at level 2. If we are talking about behavioral change, we are really talking about Level 3 and above on the SANS maturity model[1] (out of five levels) [1]. Phishing awareness training platforms are not built to give us the insight needed for those higher SANS levels because they are built for their most common use case, which is compliance-based training. Compliance-based training is centered on activity, so the metrics are built around quiz completions, module completions, training completions, and things of that nature. There is an inflection point where those activity metrics stop providing meaningful guidance in the higher SANS levels. Activity metrics are not a proxy for impact beyond compliance, particularly when they are not measured from organizational baselines that help isolate training impact. If we want to move into these more mature levels, we are going to have to evolve beyond the current measurement framework. It is difficult to credibly use a single activity-based measurement framework across all five SANS levels.

Dr. Chris Fennell: In our company we have internal training, and one of the challenges is how do we decrease people from clicking on things? We have a continual non-punitive phishing training program; we don't punish people for clicking during training, which I believe is common. We've seen similar effect sizes to what was found in Grant's research study. We vary our training with different levels of difficulty. But I think Grant is right that, even with cybersecurity training and phishing education, you only get a 2 to 3% benefit from embedded phishing awareness training programs. However, if that small benefit translates to a reduction in the burden to our SOC (security operations center) by even half of a percent, at our scale with 2.1 million associates (employees) worldwide, that's a great justification for a continual embedded phishing awareness training program.

Published research shows the game has sort of changed. Security control layers and automated tools around phishing have gotten so much better that they might have reduced the efficiency of phishing training programs. That 2 or 3% was really important when the click rate was much higher. But now, the volume of phishing emails that are actually getting through is much lower, so I wonder if there is some evidence to suggest that the conclusions in Grant's paper are right. Yet at the same time, I still feel like increasing the overall security culture, and therefore reducing the burden to the SOC by even just a quarter percent, translates to a large reduction in work hours dedicated to addressing phishing attacks.

Grant's research study also looked at the effects of static training (clickers are directed to an educational webpage) versus interactive training (clickers are directed to a webpage that provides an example phishing email and walks them through a hands-on training exercise). What are your thoughts about the difference between the type of interventions that are provided to people in phishing awareness training?

Ho: There's two answers, one with more nuance. Our study showed that, at a high level, the type of intervention did not make a difference because, when people get simulations, they don't

[1] Organizations use the SysAdmin, Audit, Network, Security Institute (SANS) Security Awareness Maturity Model to benchmark the maturity of their program and as a strategic roadmap to both plan and communicate the impact of their program [2].

interact with the training at a very high rate. Most people click out of the training, and there's nothing we can do to force them to do it short of penalties; we've seen this at the institution in the study and a couple other institutions. On the other hand, we did see that some people were willing to voluntarily complete the training. The more nuanced response is that we did see a big difference between the people who took static training and the people who did the interactive training in which they were asked questions tailored to the simulated phish they received. Static training showed no benefit. There was an improvement for those who did the interactive training. The people who did the interactive training were 20% less likely to fail future simulations than people who didn't do the training. Part of that may be because interactive training does provide better benefits for the people who are willing to learn. It's just that in our study, a lot of people didn't want to do the training voluntarily.

Johnson: It depends on the industry. My work tends to fit with higher-risk industries that strive for security beyond compliance, such as the military, where I do readiness exercises. Based on my experience, when people in these industries fail a phishing test, they take it seriously and are not going to just ignore it. In these environments, operational knowledge of how much damage a single incident can cause is built into the security culture. So they are not just going to open and close a landing page, or at least not simply return to it later. I agree with what Grant and Chris have said so far, but the people and organizations I work with do interact with training a little differently.

Fennell: One thing that I've seen, being a researcher in industry, is that a lot of phishing papers fail to cover user context and the mediating effect of if clicking links is a part of your job (for research on how user context influences click behaviors, see [5][6]). If you have 2,000 emails and 90% of them are part of your job, how do you determine what links are OK to click on and what aren't, especially if the phish looks like part of the requirements for your work? Some of my research I presented at SparkCon was trying to disentangle that [7]. Why punish perennial clickers if their entire job is looking at 2,000 emails a day, and they all have links that people have to click on? How can you punish them? I would look at it based on my company–we do business with everybody, so everybody is external, so how are employees supposed to disentangle that? This also puts an emphasis on extra security controls. So we train them; they clearly have been trained. They click, but they struggle with clicking because of their role.

If click rates don't reflect true resilience, what should we be measuring instead to determine the effectiveness of phishing awareness and training efforts?

Johnson: I think the question we have to ask is why don't click rates reflect true resilience? The problem we have is that platforms attempt to measure resilience using benchmarks where they're taking the click rates of each campaign and then measuring you against external organizations. But they're not accounting for confounders, so the benchmarks are really only a series of snapshots. So when you take a bunch of benchmarks and plot them on an X-Y axis, it implies that there's a trend and that benchmarks can speak to each other to tell a story over time. However, that becomes misleading because it conflates benchmarks with baselines. Now, if you were to take click rates and build true baselines using scientific rigor, you can create a measurement methodology that actually tracks impact from within the organization itself. Then you can track impact longitudinally and finally have the insight required to start driving improvements within your own program. So the problem is that click rates don't reflect true

resilience because they've never been measured properly to serve that purpose for the organization.

Human risk management (HRM) is the latest evolution of this same issue, reframed as a quantifiable risk score that claims to measure impact over time. It is largely the same problem presented in a different way because people believe that HRM platforms or AI (artificial intelligence) can create this magical risk score, but there are no objective measures within the risk score calculations themselves that measure impact from baselines within the organization. In other words, a quantified score does not solve the underlying measurement problem if the score is still built on non-baseline inputs. That's where we need to evolve in terms of measurement.

Ho: I like the idea that we should be measuring things in a scientific way. In our study, we measure things with randomized control trials. There's a joke saying if you really want to reduce your click rates across the board, you give people a very hard phishing email in month one and a very easy phishing email in month two, and boom, your click rate drops. But that's not how you want to do things. You want to do things in a controlled, scientific manner.

We didn't really focus on this in our study, but there is a belief among people that simulated phishing can be used in a positive way to change people's behavior to not just avoid clicking, but to increase the reporting of suspicious activity and phishing attacks. To Chris's point, we need to increase people's ability to report phishing emails accurately in a way that doesn't increase false positives and dramatically exploit false alarms. You don't want people to not click on anything, right? There's a model where you fire people for failing phishing attacks, but then you get no one answering email. Promoting proactive behavior, like accurate reporting of phishing emails, is a way to gain some benefit of behavioral change and culture change, increasing how much people are willing to interact with the security team in reporting.

Fennell: I agree with Grant. One of the things that I have seen, anecdotally, is that if you get people to report, they're going to report a lot more things and you get the false positives. It turns out false positives in nature are actually really important. Like if you see a tiger in nature and you run away, that's a good thing. Speaking for myself, we would rather have the state that we over-report false positives than under-reporting because people typically don't necessarily understand if something is really a legitimate threat or not. The under-report is when people say, "Oh, that's not a phish," and it's a false negative, and then they don't tell the security team at all. That's a real problem – that they didn't actually report the phish, they interacted with the phish, and then the security team doesn't know about it. And so it's kind of an interesting aspect of "What should we measure?" There are different ways to capture the actual reporting metric. Some of my earlier research showed that it's not just phishing that's the problem [7]. It's the usability of the tool itself which makes it difficult for people to report phishing emails. I agree with Grant and Cary; I think there are other measures that are better than click rates, and one of them is how to report.

How can organizations promote a broader culture of phishing awareness?

Fennell: I had the benefit after finishing my PhD to come to a program that was already well established, and the human-centered cybersecurity team here had spent years creating a good security culture. Some of my early interviews in the phishing realm were just asking people, "How did you learn about phishing training, and what did you learn?" Some of the evidence was

that people were saying, “I've worked at a lot of companies before, and then when I came here, I didn't realize how serious the entire company took security training.”

Our information security team took a lot of time curating what we wanted our associates to know. Some are common things like phishing. Our associates have multimodal ways to access phishing training beyond just reading the training email, like clever videos that have been created about phishing. They created narratives around the training. We’ve seen some interesting research [8][9]. I think we’ve taken those things and now have built an entire training program that just creates a fantastic culture of awareness. So when people come here, it’s not common to think, “Oh, I just click this link;” a lot of people are very skeptical about suspicious links and they pay attention to external markers like where URLs (uniform resource locators) redirect.

Johnson: I know it's a touchy subject, but in my experience across organizations where phishing simulation programs are built to benefit the organization, it can create a broader culture of phishing awareness if you're creating baselines within the organization for both click rate and report rate and then a ratio thereof. In my experience, phishing simulations tend to be more important than many platforms acknowledge. You can train people and give them all sorts of content, and they can ignore it and not absorb it, but they have no choice when they get a simulation. They have to react one way or another. That's the one way to give them that visceral reaction that they will remember. Obviously it's a combination of two types of training – training content and doing phishing simulations – on a regular cadence.

Ho: The term “awareness” gets thrown around a lot, but the number one message of our study, I think if anything else, is broader: organizations should think carefully and concretely about what they're trying to achieve. What does “awareness” mean? How does that translate into something that we can actually measure as an organization? But if you read NIST 800-53 for example, there are a lot of recommendations that organizations should implement for different security practices [10]. It's a matter of prioritizing time, practicing budget, and employee attention. As long as organizations are clearly defining what they're trying to get out of awareness and measuring that in a rigorous way, they can determine value. That's the most important thing we want to make clear out of the research study.

Fennell: We're talking about the top 10 Fortune companies. What happens to the Fortune 500 companies where they have one security guy? He's in charge of the SIEM (Security Information and Event Management) and also has to try to do a phishing training program. He doesn't have the benefit of the same resources, so he's learning from those policy places, or NIST or other places to try to implement those things. It's even more finite resources. That's really where we should spend time, this middle layer of enterprise, 1,000 employees to 10,000 employees that I would consider a small or medium business.

Where is the line for "realistic" testing? Is it acceptable to use real-world lures, or might those cause damage to the employer-employee relationship?

Ho: In our study, the organization’s leadership was very concerned about using COVID as a topic, “You're fired,” etc. I appreciate the real-world challenges here, but at the same time, what are the goals? If organizations can clearly articulate that the goal does not involve real-world use, that's fine. But if the goal is to simulate realistic phishing attacks to train employees for realistic phishing scenarios, then I think you have to use realistic phishing. If there's a way in which

you're viewing awareness as not trying to train employees against real phishing attacks, realistic lures aren't the objective. If the goal is broader cultural change that you're trying to measure some other way, you probably don't need them. But if you want to train people to spot a phish, if you want to give employees exposure to real phishing attacks, attackers are not going to say, "Oh, I think this topic might make them upset. I'm not going to use it." It depends on your goal. If you're trying to train them for something, you should probably train them in the most accurate way possible. There is this tradeoff though.

Johnson: This is a tough one because there are so many factors that go into this. In certain highly regulated industries or high-risk environments like the military, it's less of an issue. With military readiness exercises, there is no issue with realistic testing. So it really depends on the industry. To give you the flip side of that, I have worked with legal and accounting firms where the political power structure is completely different, where you have lawyers and accountants on equal footing with the cybersecurity team. As soon as you cross that line, you're going to find out, and they're going to point the finger at IT security. Within the industries I work with, this is typically not an issue, however, there are always some lures that just cross the line regardless. You can't talk about bonuses, especially at Christmas time. You can't talk about anything that relates to firing. Sometimes platforms can cross the line by even having some of these topics in their phishing simulation libraries because practitioners using these platforms assume that the content is vetted. That's where you get these situations where someone uses a bonus simulation at Christmas time, and it ends up in the news. A lot of factors go into that, but it does depend on the industry.

Fennell: What Grant and Cary said really resonates. Organizations obviously cannot do some real-world lures. They just can't from a legal perspective, like, "Hey, click on this email, you'll get harmed," or some other effect. We just can't do it partly due to institutional review boards and legal controls. That being said, maybe the real question is, can we do something where we've trained them to be skeptical and understanding of the semantic training? That way, when they do get real-world phishing emails, the training will translate. I think we can have realistic enough training that when people do get real-world phishing emails, they get reported.

Johnson: Additionally, you cross the line when you've taken all the clues out. If you're running simulations and there's no clues in it, there's no educational value. We can't expect attackers to stop phishing using emails without any red flags in it. Touching on Chris's point there, that's extremely important.

In a cybersecurity context, psychological safety is the belief that one will not be punished or humiliated for making a mistake or raising a concern. How do we maintain psychological safety within an organization when these phishing simulations are, in essence, a form of deception?

Johnson: We need to understand why it's perceived as punishment. Traditionally, phishing simulations have always been, and still are, used to feed a training platform content engine. Phishing emails are pushed out, the employee fails, and then they get content, training, and, in some cases, more simulations. Of course, that is perceived as punishment.

If you were to flip the use of simulation on its head and run simulations in such a way that you are giving people a roadmap, saying, "We're going to give you simulations over a period of time,

say six months to a year, and then we're going to accumulate the data to do a Pareto analysis and determine who are the never clickers, who are learning from simulations alone, and who are the movers that continue to struggle—the 10 to 15% creating 80% of the clicks," you're giving them plenty of visceral experiences without punishment. Then you can identify who really needs training. When you approach it that way, and you go to that 10 to 15%, it's still about perception. You want to tell them, "Look, you have proven beyond a doubt that this is a problem, and we need to help protect you both at work and at home in your personal life. You need the street smarts to protect yourself outside of work just as much as at work." So when you think about using simulations and data for the benefit of both the organization and the employee, you can start to eliminate some of the issues that can build a culture of resentment.

Ho: Do you do negative or positive rewards for phishing? I think there traditionally has been a lot of thought that we should be penalizing people for failing, and I, for a number of reasons, don't think that–for example, firing people for failing at simulations–is necessarily the best idea. To Chris's earlier point, some people's roles depend on doing this work. It can also be very hard to detect phishing attacks, especially with AI. I think, based on data I've seen, based on what I think of in terms of security, penalizing users, firing them, threatening to fire them is one way to not do that.

I have seen people incentivized for reporting. If you're doing well at reporting, if you're reporting a lot of simulations, you get bonuses. That can be a way to benefit culture.

Fennell: To add onto that, I think if it's punitive, that doesn't really lead to behavior change. The behavior change is rewarding the positive actions, the reporting. That's sort of the approach that our internal security team adopted. By doing that, we incentivize people to say, "Oh yeah, I did the right thing. I reported the phish, the training simulation– boom!" You can add gamification layers to it–it seems like a cheesy thing to have–but people treat it really importantly around here. I've met a lot of people across our Walmart Global Tech area that have said, "Oh yeah, I saw the simulation. I can always catch it." And they're very proud of that, which is good. That means we're speaking to the culture that we want them to have, and so we're incentivizing them to say, "Hey, it's non-punitive." We're incentivizing associates to report. I think that's the right positive action.

Threat actors can now use Generative AI to eliminate the classic warning signs in phishing emails. How must phishing awareness training evolve to address this concern?

Ho: We currently have a separate, in-progress study looking at AI-generated phishing and trying to measure the risk and where that comes from. But, even when we were designing the training for the phishing training effectiveness study, it was very difficult to identify an easy set of tips that average people can use on a regular day-to-day, every-single-email basis. So this idea – that the goal is to avoid people clicking on links – is going to be very difficult to do in this current form. If we're trying to train people to do that, we probably need better support tools that can help notice things, for example, building new systems that can help give even simple warnings like "Hey, this thing is external."

Training would probably need to co-evolve with the design of tools that assist people for spotting these kinds of things. It's not realistic to tell someone, "You should check every single header of every single email that you get." For a lot of these vendor-related platforms, what they'll do as a

service is rewrite the URLs to a vendor domain, which allows them to basically block it later. What that means is, if I hover over the link and I try to look at it, it's always a vendor URL. So, I can't use a lot of the tips that we've traditionally given people. I think training would probably need to evolve with tools that we'll need to build to help people to identify these kinds of attacks.

Johnson: When it comes to AI phishing, a lot of the clues have been removed. So, you don't want people looking around for spelling and grammatical errors, and their eyes are shifting everywhere. You have to have them focusing on where the clues tend to remain, no matter how sophisticated the attack. That tends to be in the sender's email address and any links. Then you need to get them to build verification skills. Verification skills are what you can apply to any type of attack. It's repeatable. It applies to any attack, whether it's a deep fake or any sort of phishing across any communication channel. The program that I developed for email phishing is called **HELP** and it basically is: **H**over over the **E**mail address and any **L**inks, then **P**ause to verify. The email address has to match where it claims to come from. The link has to have a domain that represents a safe and expected internet destination within the context of the actual email. Then pause to verify. Don't get caught on any step, just move to the next one if you're not certain because, ultimately, pausing to verify through another out-of-band channel is the one activity that is most effective across any attack-type.

With the rise of AI voice cloning and SMS[2]-based attacks, how do we train a workforce that is increasingly "mobile-first" and susceptible to social engineering outside of the traditional email inbox?

Johnson: I'll continue with the same vein. It's an opportunity to create positive change with simplicity. I'm noticing now there's even some deepfake-as-a-service kits on the black market. As AI attacks become more pervasive, they're going to end up doing a lot of the awareness work in and of themselves. We need to get people to start building verification skills, that is, a simple pause, verify, continue. It applies across all types of attacks. As this all gets more common, more sophisticated, more pervasive, we need to backtrack, simplify, and focus on verification skills-just operationalize verification skills across any channel.

Fennell: There's two ways I can see approaching this. One is the technology approach. As technologists, we can continue developing tools to make it easier for people to understand. But then from the other side, from the sociotechnical side, it's a semantic problem. How do people assign meaning to text, or how are they listening and understanding those things? Those things we can still train them for. We can say, “Is the contextual plausibility of what I'm being asked to do something that I should be doing?” Introducing skepticism in that space is one thing.

In terms of SMS phishing, it's just going to get better and better. So assuming that SMS, instead of having broad phishing, it's all going to be spear phishing, it's all going to be targeted and the voice or even the text and the mannerisms that are being used are near perfect, we still have to ask, “Is this thing what I should be doing, this thing what I'm being told to do?” That semantic check is the last social barrier that we will always have. It's like behavior change. It's nothing new. We've done this before. It's just the mode has changed and maybe even the velocity of what AI allows us to do is even faster than before. I think the first spear phishing research paper I was

[2] Short Message Service like text messaging on mobile devices.

just looking at was from 2004 [11]. So this is not new. We've been dealing with this problem for a while.

Johnson: You have to teach people that "Spidey sense." That's the biggest challenge. Even myself, I live and breathe phishing, and on occasion, I'll start to react to something before thinking sometimes.

Ho: I agree that verification is going to be the most important thing. With things like vishing and deepfakes, it's actually important for organizations to build into the process of sensitive actions some sort of operational verification through some safe out-of-band channel that should be part of the institutional process, so it isn't easy for someone to just do something without that verification. Then teaching people to follow that verification process will be the way going forward.

Fennell: I would push back only a little bit in the fact that at the scale of our organization, adding verification is absolutely essential, but the pace at which we do business often outpaces our ability to add the verification.

Johnson: You have to have the "Spidey sense" to know when to verify. Otherwise, if you verify everything, the business comes to a halt.

Fennell: Our security research team talked about how it would be great if we could do formal verifications as part of the process, part of the development before it even gets to the fact that we have those security layers and controls. I think that's that tension that we'll always have. "Hey, here are these new LLM (large language model) technologies. We should be using these everywhere." What does that mean? It means "I'll be able to automatically process all these emails." How are they not susceptible to phishing since they were literally designed around these tools? So I can see the benefits of what you're saying. But what happens when those verifications are replaced by agents and now we are using a probabilistic system to handle things that should be deterministic? That's a really scary thing. Having worked in a SOC – because my team is placed within a SOC – I can tell you that the challenge is very real and very difficult.

Ho: I appreciate the realities of the business operations. To be honest, this is where we are going to experience some pain in the short term. It will have to become a cost of business that either you're going to accept falling victim to this or you're willing to build in this friction of verification. Because ultimately we're going to get to a point where I don't think everyone will be able to have that "Spidey sense" because things will be too difficult to do. It's going to be this business calculus that will shape up.

Johnson: That's where it varies by industry. When you get high-risk industries, they're going to be more willing to operationalize verification skills, whereas the ones more focused on compliance, they're just going to refuse that. They'll never react in that way.

We've been talking about phishing training, which is often a part of a broader cybersecurity awareness and training program within organizations. If you could change one thing about how organizations approach cybersecurity awareness today, what would it be?

Fennell: The security controls around training people to have better security practices are so important. That doesn't even have to limit itself to phishing. We can think about secure coding, developing with security first or formal methods. If you embed the culture, it saves you so much time later on in terms of fixing it, remediation, reducing the time to remediate. Because developing is the cheap part. If we can develop it with a security culture first, that's fantastic. That's the never-ending battle. Even here with such a good culture, we're still fighting that battle to get people to think security first, because security is a second-order concern. Making that a first-order concern, or at least a part of the conversation, allows us to save so much heartache later on.

Johnson: We need to backtrack. Go back to the foundation and start from the very beginning. It's difficult to improve what you don't measure properly. We need to create baselines within the organization itself so that we know what is and what isn't working for our own environment. Benchmarks are comparing our results to external organizations that aren't even measuring the same thing for the same reasons, so this vertical type of siloed metrics don't help us. We need to start with the proper foundation. We need to go back to the basics and start working on how we are measuring what works within each different environment. That requires baselines, both click rate and report rate baselines, with the support of some of the subjective metrics that we create. We need that North Star measurement framework, something that measures impact objectively to guide optimal programs. Why not use phishing results properly for this? It's the most common, and arguably the biggest human risk. That's where we need to start.

Ho: The core message of our work is, if you believe any security control that you have is adding value–phishing training in particular–then you should be able to clearly articulate what that value is and hopefully be able to scientifically measure that. That takes some care. There are all sorts of ways to measure phishing training in flawed ways. But really scientifically measuring what is the value we think we're getting, and going out and measuring if we're getting that value is so important. That was sort of the goal of our study and the focus of our work. These training platforms are expensive. They cost millions of dollars, potentially. So if you're spending a lot of money on this, you should be able to identify what the value is and measure it right. That's so important.

Disclaimer

Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and roundtable experts and do not necessarily reflect the views of their employers, the National Institute of Standards and Technology, or the U.S. Government. Certain commercial companies or products are identified to foster understanding, not to imply recommendation or endorsement by the National Institute of Standards and Technology, nor to imply that these are necessarily the best available for the purpose.

Disclosure

Google Gemini 3 was used to refine the language of questions asked during the roundtable, in accordance with the authors' instructions. All generated content was reviewed, edited, and verified by the authors to ensure accuracy and originality.

References

1. SANS Institute. (2025). Embedding a Strong Security Culture: SANS 2025 Security Awareness Report. https://www.sans.org/mlp/ssa-security-awareness-report
2. Spitzner, L. (2021). Strategically managing your human risk – Leverage the Security Awareness Maturity Model. SANS Institute Blog. https://www.sans.org/blog/strategically-managing-your-human-risk-leverage-the-security-awareness-maturity-model
3. Ho, G., Mirian, A., Luo, E., Tong, K., Lee, E., Liu, L., Longhurst, C. A., Dameff, C., Savage, S., & Voelker, G. M. (2025). Understanding the efficacy of phishing training in practice. In *2025 IEEE Symposium on Security and Privacy (SP)* (pp. 37–54). IEEE. https://doi.org/10.1109/SP54013.2025.00012
4. Lain, D., Kostiainen, K., & Čapkun, S. (2022). Phishing in organizations: Findings from a large-scale and long-term study. *2022 IEEE Symposium on Security and Privacy (SP)*, 842–859. https://doi.org/10.1109/SP46214.2022.9833766
5. Greene, K., Steves, M., Theofanos, M., & Kostick, J. (2018). User context: An explanatory variable in phishing susceptibility. *Proceedings 2018 Workshop on Usable Security*. https://doi.org/10.14722/usec.2018.23016
6. Dawkins, S., & Jacobs, J. (2023). *NIST Phish Scale User Guide* (NIST Technical Note 2276). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.TN.2276
7. Fennell, C. (2023). *Email Challenges and Phishing with Chris Fennell*. Sp4rkCon 2023. https://internal.walmart.com/content/sp4rkcon/wmeo/sp4rkcon-home/sp4rkcon-2023/agenda/email-challenges-and-phishing.html
8. Kumaraguru, P., Sheng, S., Acquisti, A., Cranor, L. F., & Hong, J. (2010). Teaching Johnny not to fall for phish: Understanding how to educate users about phishing. *ACM Transactions on Internet Technology*, 10(2), 1–31. https://doi.org/10.1145/1804601.1804607
9. Rader, E., Wash, R., & Brooks, B. (2012). Stories as informal lessons about security. *Proceedings of the Eighth Symposium on Usable Privacy and Security*, 1–17. https://doi.org/10.1145/2335356.2335364
10. National Institute of Standards and Technology. (2020). *Security and privacy controls for information systems and organizations* (NIST Special Publication 800-53, Revision 5). U.S. Department of Commerce. https://doi.org/10.6028/NIST.SP.800-53r5
11. Caputo, D. D., Pfleeger, S. L., Freeman, J. D., & Johnson, M. E. (2014). Going spear phishing: Exploring embedded training and awareness. *IEEE Security & Privacy*, *12*(1), 28–38. https://doi.org/10.1109/MSP.2013.106

Author Bios

Shanée Dawkins is a computer scientist at the National Institute of Standards and Technology, Gaithersburg, MD 20899, USA, where she leads the Phishing project in the Human-Centered Cybersecurity program. Dawkins received a Ph.D. in computer science from Auburn University. Her research interests include human-centered design and usability, with a particular emphasis on developing standards and guidelines that enhance the user experience in complex systems. Contact her at shanee.dawkins@nist.gov.

Julie Haney is a computer scientist at the National Institute of Standards and Technology, Gaithersburg, MD 20899, USA, where she leads the Human-Centered Cybersecurity program. Haney received a Ph.D. in human-centered computing from the University of Maryland, Baltimore County. Her research interests include the work practices of cybersecurity professionals, the implementation of human-centered cybersecurity approaches into practice, and people's cybersecurity perceptions and experiences. Contact her at julie.haney@nist.gov.

Published in IEEE Security & Privacy Magazine July/August 2026 issue