Gone phishin': CyLab exposes how our ability to spot phishing emails is far from perfect

October 3 2016

Each year, tens of millions of phishing emails make it to your inbox, uncaught by your email client's spam filter. Of those, millions more slide past our own judgment and are clicked and opened. A recent study out of Carnegie Mellon's CyLab Security and Privacy Institute has revealed just how likely we are to take the bait.

"Despite the fact that people were generally cautious, their ability to detect phishing emails was poor enough to jeopardize computer systems," says Casey Canfield, a CyLab researcher from Carnegie Mellon's Department of Engineering and Public Policy.

Canfield's study was recently published in the journal *Human Factors*. Those interested can test their own phishing email detection skills in our brief online quiz.

In the study, Canfield and her colleagues showed a set of participants information about phishing before asking them to evaluate 38 different emails, half of which were legitimate and half were phishing. For each email, participants answered questions about whether the email was phishing, what action they would perform, their confidence in their choices, and the perceived consequences of falling for the email if it was phishing.

On average, participants were only able to correctly identify just over

half of the phishing emails presented to them. Fortunately, participants displayed a little more caution when it came to their behavior: roughly three-quarters of the phishing links were left un-clicked.

"Some users were able to identify a vast majority of the phishing emails, but only because they were biased to think everything was a phishing attack," Canfield says. "So they didn't necessarily have a high ability to tell the difference between phishing and legitimate emails."

What's more, participants' confidence levels were not always calibrated with their ability.

"When making decisions about phishing emails, people were more cautious when they were unconfident and perceived very negative consequences of opening a phishing email," Canfield says. "Unfortunately, they were often overconfident so they would still fall for phishing attacks."

Based on the results, the authors of the study suggest interventions such as providing users with feedback on their abilities and emphasizing the consequences of phishing attacks. One effective training method that companies commonly use, Canfield explains, is sending out fake phishing emails and teaching a user about phishing emails if they open the email. This training method, called "embedded training," was originally developed by the CyLab Usable Privacy and Security Lab.

"It seems like those trainings may not always be making people better at telling the difference, but it's probably making them more cautious," Canfield says. "Helping people tell the difference may not be as useful as just encouraging them to be more cautious."

Other authors on the phishing study included Engineering and Public Policy professors Baruch Fischhoff and Alex Davis.

More information: C. I. Canfield et al, Quantifying Phishing Susceptibility for Detection and Behavior Decisions, *Human Factors: The Journal of the Human Factors and Ergonomics Society* (2016). DOI: 10.1177/0018720816665025