The Big Phish: Cyberattacks Against U.S. Healthcare Systems

Adam Wright, Ph.D.[1,2,3], *Skye Aaron, B.A.*[1], *and David W. Bates, M.D., M.Sc.*[1,2,3]

[1]Brigham and Women's Hospital, Boston, MA, USA; [2]Harvard Medical School, Boston, MA, USA; [3]Partners HealthCare, Boston, MA, USA.

KEY WORDS: security; privacy; phishing; electronic health records.
J Gen Intern Med 31(10):1115–8
DOI: 10.1007/s11606-016-3741-z

From: [user]
Sent: Tuesday, November 11, 2014 11:52 AM
To: [user]
Subject: Faculty And Staff Mailbox Alert.
Your password Will Expire In The Next TWO {2} Days Current Faculty and Staff Please Log On To IT WEBSITE To Validate Your E-mail Address And Password, Or Your E-mail Address Will Be Deactivated.Thank You.
ITS help desk
ADMIN TEAM

©Copyright 2014 Microsoft
All Right Reserved.

Phishing, the practice of obtaining computer credentials from users through manipulation or deceit, dates back at least 20 years to America Online (AOL), where users would impersonate AOL staff members and send instant messages to other users convincing them to disclose their passwords or credit card numbers. The term itself was coined by Koceilah Rekouche, a hacker known online by the pseudonym "Da Chronic," who created a tool for automating and accelerating this process in 1995.[1] The manual process had sometimes been called fishing (as in fishing for passwords), and Rekouche termed the password-stealing function of his software "phishing"—the term stuck, and the behavior has subsequently expanded far beyond AOL over the last two decades.

The email above was sent to users at our hospital and is one of many like this we receive every month. It encourages recipients to click a link where they are asked to enter their username and password. However, the site is operated not by our IT department, but by hackers seeking to gather passwords. When a user takes the bait and enters a password on the hacker's site, the hacker gains the ability to access a range of online services by impersonating the user. While most users who receive an email like this one should know better than to click the link, phishing exercise results show otherwise. Users do fall victim to these manipulations, and some provide information, such as passwords, that is useful to hackers.

The success of phishing messages is often tied to realism and authority—they may appear to be from an authority such as a hospital IT department and warn users that their accounts will be shut off if they don't "update" them by entering their passwords. Phishing websites, which users access after clicking links in emails, are often designed to mimic institutional sites with misappropriated logos and similar designs, and they have addresses that resemble official sites, sometimes with minor misspellings or a lowercase letter L replaced with the number 1. Over time, phishing attacks have become more sophisticated, with higher quality emails and more convincing sites for capturing credentials.

Although many phishing attacks are indiscriminate, targeting large numbers of users, a variant called "spear phishing" focuses on smaller groups of users or even specific individuals. Spear phishing attacks can be particularly effective because they can be carefully targeted to the sorts of links and deception most likely to trap a particular user—for example, a note apparently from the user's boss or even a journal that the user regularly submits to.

PHISHING ATTACKS AGAINST HOSPITALS

Phishing attacks like the one above are widespread, and organizations in most industries, including healthcare, have fallen victim to them. From press accounts and public announcements, we identified at least ten incidents since 2014 where hackers gained unauthorized access to hospital systems through phishing in the United States (Table 1), including two separate attacks against our organization. We believe that this list is almost certainly incomplete, as the vast majority of phishing attacks go unnoticed or unannounced, and some security consultants have reported that hospitals routinely undergo several phishing attacks every week. Although organizations are required to announce breaches of protected health information (PHI), not all phishing attacks lead to disclosures of PHI, nor are all investigated.

CONSEQUENCES OF PHISHING ATTACKS

Once credentials are stolen by hackers, they can be put to a variety of uses. Most commonly, hackers use them to misappropriate identifiable information—often by searching the compromised user's mailbox for spreadsheets or other documents

Published online May 13, 2016

Table 1. Phishing Attacks Against Healthcare Systems in the United States

Institution	Approx. Month of Attack	Patients Affected	Data Potentially Disclosed
Baylor Regional Medical Center, Plano, TX	Jan 2014	1981	"patient information, including names, addresses, dates of birth, or telephone numbers, some clinical information such as treating physician, department, diagnosis, treatment received, medical record number, medications, medical service code or health insurance information and Social Security numbers" (Source: http://www.databreaches.net/tx-baylor-regional-medical-center-at-plano-notifies-patients-after-physicians-fall-for-phishing-attempt/)
Catholic Health Initiative Franciscan Medical Group, Tacoma, WA	Jan 2014	8300	"demographic information (for example, name, address, date of birth, telephone number), clinical information (for example, treating physician and/or department, diagnosis, treatment received, medical record number, medical service code, health insurance information), and in a small number of instances, Social Security numbers" (Source: http://www.chifranciscan.org/news/Notice-to-our-Patients-Regarding-Phishing-Scam/)
Partners HealthCare System, Boston, MA	Nov 2014	3300	"patient demographic information, such as names, addresses, dates of birth, telephone numbers, and, in some instances, Social Security numbers, and some of our patients' clinical information, such as diagnosis, treatment received, medical record numbers, medical diagnosis codes, or health insurance information" (Source:https://web.archive.org/web/20150719014105/http://www.partners.org/privacy-incident-notice/)
Seton Family of Hospitals, Austin, TX	Dec 2014	39,000	"demographic information (i.e., name, address, gender, date of birth, etc.), medical record numbers, insurance information, limited clinical information and, in some cases, Social Security numbers [but not] individual medical records or billing records" (Source: https://web.archive.org/web/20150717053211/http://www.seton.net/email_phishing_incident_at_seton_family_of_hospitals)
St. Vincent Medical Group, Indianapolis, Indiana	Dec 2014	760	"patient's name, demographic information such as date of birth and phone number, account numbers, limited clinical information related to services the patient had received and, in some cases, social security numbers [but not] individual medical records or billing records" (Source: http://www.stvincent.org/uploadedFiles/SV_Health/St.Vincent_Medical_Group/HRKHL-1731388-v2-SVMG_Substitute_Notice_Email_Phishing_Hacking_Attack.pdf)
Middlesex Hospital, Middletown, CT	Oct 2015	946	"name, address, date of birth, medical record number, medication, date of service, and/or diagnosis" (Source: http://middlesexhospital.org/press-releases/e-mail-phishing-scam-results-in-data-breach)
Brigham and Women's Hospital, Boston, MA	Nov 2015	1009	"name, medical record number, date of birth, date of service, provider name, health diagnoses and treatment information" (Source: http://www.hipaajournal.com/phishing-attack-suffered-by-brigham-and-womens-hospital-8272/)
Oakland Family Services, Pontiac, MI	Jul 2015	16,107	"names, internal client ID numbers, dates of service and types of service provided," and in limited instances, "dates of birth, telephone numbers, addresses, diagnoses, health plan ID numbers, insurance numbers and social security numbers" (Source: http://www.oaklandfamilyservices.org/OaklandFamilyServices9-10-15.pdf)
Metropolitan Jewish Health System, Brooklyn, NY	Jan 2016	2483	"member and patient names, member numbers, diagnoses, treatment dates, and the facility where members were recently treated" (Source: https://www.mjhs.org/privacy-statement/notice-regarding-phishing-email-incident/)
City of Hope Hospital, Duarte, CA	Jan 2016	1024	"elements of protected health information, such as patient names, medical record numbers, dates of birth, addresses, email addresses, telephone numbers and some clinical information such as diagnoses, test results and dates of service"; fortunately, "only one Social Security number was exposed" (Source: http://www.cityofhope.org/news/city-of-hope-responds-to-phishing-email-attack)

that contain personally identifying information for staff or patients. This identifiable information is most valuable to hackers when it contains sensitive identifiers like dates of birth and social security numbers. These stolen identities can be sold through online black markets or local criminal networks and used for a range of financial fraud, including filing false tax returns or applying for credit. Identities stolen from healthcare providers often fetch much more than even stolen credit card numbers—a recent black market listing profiled in the news media offered a "value pack that includes ten people's Medicare numbers" for $4,700,[2] although other experts put the value of stolen health information much lower, at around $10 per identity (still multiples of the value of a credit card number).[3]

Stolen network credentials can also be used for other types of fraud. For example, hackers sometimes use stolen credentials to access payroll systems and change salary direct deposit destinations to bank accounts they control, allowing them to steal wages.[4] In our review, we did not uncover cases where attackers were making healthcare-specific use of stolen credentials, such as forging prescriptions or stealing clinical data to use for blackmail or other nefarious purposes. However, there is a risk that hackers may increase their sophistication or explore these other options, particularly as the financial industry increases its safeguards.

Hackers also routinely use stolen credentials to launch further phishing attacks, part of why phishing attempts sometimes appear to come from someone known to the user. Hackers particularly seek credentials with elevated privileges, such as those belonging to network administrators, as such users often have access to create new accounts, modify account privileges (e.g., grant additional privileges to other users), or access databases and fileservers directly, bypassing normal security and monitoring measures.

Once credentials have been obtained and the network breached, hackers can set up a "beachhead" to launch other attacks—for example, installing malware. Several recent cases of ransomware attacks have affected hospitals in recent weeks, with Hollywood Presbyterian Hospital experiencing extended down times before eventually paying hackers a $17,000 ransom. Although it appears that hackers used other vehicles besides phishing to perpetrate the Hollywood attack, phished credentials are a very effective approach for mounting more sophisticated malware attacks.

In addition to these direct consequences, organizations that fall victim to phishing may suffer reputational harm and may also have to bear additional costs, including fines from regulators, damages paid to patients for direct harms, as well as the cost of providing credit monitoring and identity theft-related services. Insurance policies are available to protect organizations from some of these risks, with this market quickly evolving, but some aspects of risk, like reputational harm, can be hard to quantify.

PREVENTING AND MANAGING PHISHING ATTACKS

Healthcare organizations and providers can employ a variety of techniques to reduce the risk of phishing attacks. The most common method organizations employ is training—teaching users to identify phishing attacks and reminding them not to respond to them. Although training may be effective for many users, it does not afford complete protection. Our organization provides regular training to our users, but some still fall victim to phishing scams, and strategies that depend entirely on user behavior are not likely to be successful, as hackers are sophisticated and persistent. We have also partnered with outside organizations to send simulated phishing emails to our users. If users click links in the email, they are taken to training sites which remind them about the nature and prevention of phishing.

Another common approach to managing phishing is filtering. The first line of filtering is detecting and blocking phishing emails before they can be delivered to users. Filtering can also be used to block access to the websites that hackers use to collect stolen credentials. Even after credentials have been stolen, filters can be used to block network connections from computer systems known to be used by hackers. Organizations can also employ filters to block exfiltration of data—for example, by automatically stopping large outbound file transfers, or screening network traffic for what appear to be social security numbers. However, as discussed above, many other attacks, such as ransomware, can be perpetrated even without exfiltration of data. Although filtering methods can be effective, they represent an arms race with hackers who continually develop techniques to evade each type of filter. As such, multiple layers of filtering, effective use of encryption, effective maintenance of software and infrastructure, malware detection, and monitoring are generally required—a strategy known as "defense in depth."

Organizations can also use techniques to limit harm once credentials are misappropriated. The most basic approach is to require frequent password changes—this limits the amount of time that misappropriated credentials can be used by hackers, at the cost of some inconvenience to users. However, while this approach limits long-term use of stolen credentials, it offers little protection against hackers who make immediate use of stolen credentials. Another approach is to limit the amount of data that users have access to—for example, by not using social security numbers as internal identifiers, or by restricting access to complete social security numbers for most users, already a best practice. In a similar vein, limiting elevated administrative privileges to only those users who need them and encrypting credentials when they traverse the network can further mitigate exposure in the event of a breach. Careful training and auditing should accompany privileges like access to sensitive information or elevated credentials.

However, the single most effective technique for mitigating the impact of phishing attacks is two-factor authentication. With two-factor authentication, users must provide another authentication factor in addition to their password, such as a biometric (e.g., a fingerprint scan) or a temporary numeric code generated by a device or application, or sent to the user through text message. With two-factor authentication, even if a hacker gains access to a user's password, the hacker won't be able to access a user's account unless he or she is also able to compromise the second authentication factor. Experience in other industries, including finance, suggests that two-factor authentication represents a significant advance for security, and is a strong protection against misuse of stolen credentials. Though effective, two-factor authentication requires users to carry or use additional devices or tokens, and it takes valuable time.

RECOMMENDATIONS FOR HOSPITALS AND HEALTHCARE PROVIDERS

Healthcare organizations should train users to recognize, avoid, and report phishing attacks, and users must exercise skepticism towards emails that pressure them into clicking a link or sharing personal information. Although security that relies on user training alone is unlikely to be sufficient, engaged and educated users is the first line of defense against many attacks. In addition, hospitals must employ multiple layers of filtering, detection, encryption, and monitoring, both to prevent breaches and to mitigate exposure in the event of a breach, and the principle of

least privilege must be applied when granting access to sensitive information and account capabilities. Finally, two-factor authentication is the most important step that users can employ to reduce the risk of harm through phishing. Although two-factor authentication requires an extra step in authentication, and may require retrofitting of security mechanisms, the security benefits are substantial, and even a prevented single attack could far outweigh any costs of implementation as well as the cost of inconvenience to users.

Acknowledgements: *We appreciate the input of Dean F. Sittig and Lipika Samal who reviewed early versions of this manuscript and provided helpful feedback.*

Corresponding Author: *Adam Wright, Ph.D.; Brigham and Women's Hospital, 1620 Tremont St., Boston, MA 02115, USA (e-mail: awright@bwh.harvard.edu).*

Compliance with Ethical Standards:

Conflict of Interest: *The authors declare they do not have a conflict of interest.*

REFERENCES

1. Rekouche K. Early phishing. arXiv preprint arXiv:11064692. 2011.
2. Shahani A. NPR: The Black Market For Stolen Health Care Data. Available at: http://www.npr.org/sections/alltechconsidered/2015/02/13/385901377/the-black-market-for-stolen-health-care-data. Accessed April 12, 2016.
3. Humer C, Finkle J. Your medical record is worth more to hackers than your credit card. 2014. Available at: http://www.reuters.com/article/2014/09/24/us-cybersecurity-hospitals-idUSKCN0HJ21I20140924. Accessed April 12, 2016.
4. Bonnett C. Employees' Direct Deposit Rerouted After Phishing Attack. 2014. Available at: https://today.duke.edu/2014/01/phishing. Accessed April 12, 2016.