AGARI

WHITE PAPER

Bolstering Trust in the Email Ecosystem

by Byron Acohido

Email is inexpensive, instant and ubiquitous. Email has become deeply entrenched as our favorite tool for sharing, collaborating, coordinating and archiving. Forget the telephone or social media, when it comes to consumers and companies communicating, email holds an unshakeable dominance.[1]

But what if email could no longer be trusted? That's not an idle, rhetorical question. Trust in the email ecosystem today is brittle at best. As a communications tool, email remains viable and dominant only due to a monumental effort to keep relentless spammers and phishers from completely undermining the trustworthiness of email messages.

Consider that enterprises spend billions each year on email filtering and malware detection technologies just to keep the daily tidal wave of spam and phishing messages from inundating their operations. Consumers, too, must pay up for anti-malware protection for their personal computing devices – and still be ever vigilant for messages carrying viral attachments or links to corrupted web pages.

Email has remained viable and dominant only because the good guys, thus far, have managed to keep the bad guys in check an acceptable amount of time. That tenuous equilibrium is always in danger of collapse because the best defenses react only to the latest form of trickery. Identity deception is at the core of what cyber criminals do as they continually vary their methodology, content to succeed to some degree with each iteration of an attack type.

The latest development in this cycle is a form of spear phishing that entirely circumvents the best filtering and malware detection technologies on the market. This paper discusses this new type of attack and examines a new approach to addressing spear phishing – one that does not rely on reactive measures and, instead, directly focuses on addressing identity deception with a durable methodology based on machine learning.

"All successful spear phishing attacks feature **identity deception** at their core."

Devastating Deception

All successful spear phishing attacks feature identity deception at their core. Criminals target specific individuals at certain companies and then craft custom messages to lure the victim into taking an action. Usually the goal is to get the victim to divulge account logons or to click on a malicious attachment or URL. Either way the attacker gets control over the victim's PC and a foothold into the organization. Spear phishing relies on manipulating humans, and busy humans will always be the weakest link. Thus a robust solution to spear phishing ought to focus on isolating, and cutting off, the sender of such messages.

This latest iteration of spear phishing is brilliant in its simplicity and has proven, in a short period of time, to be devastatingly effective. Here's how it works: The attacker sends a senior executive an email impersonating someone with whom the executive has a close relationship. Very simply, the executive falls for it – the email sender convinces the victim to forward sensitive documents or to take steps to execute a cash transfer into an account the sender controls.

This latest type of attack relies on a deeply researched, meticulously crafted and very persuasive message sent specifically to the victim at an opportune moment. The instigating email is a one-off message, not part of a mass-mailing; there is nothing in the header to set off any alarms; and it carries no malware-laden attachment, nor any malicious URLs. In short, there is nothing for message filtering and malware detectors to react to. It is an email-enabled breach mechanism that relies 100% on social engineering.

Law enforcement refers to this new type of spear phishing as Business Email Compromise, or BEC scams. It has also been referred to as "whaling," "human hacking" or "CEO fraud." BEC scams plagued an unprecedented number of businesses in 2015. According to the FBI, a surge of BEC capers resulted in scammers stealing a stunning $1.2 billion from more than 7,000 U.S. companies[2] from October 2013 through August 2015. Media reports suggest these attacks have not slowed down, and logic dictates that they will persist until such time as executives stop falling for such ruses.

Since BEC spear phishing cannot be blocked by legacy email filtering and malware detection technologies, the obvious way for executives to stay out of the line of fire might seem to be to stop trusting email. In fact, it is a safe assumption that distrust of email has already started taking hold, as word circulates about cases such as Ubiquiti Networks losing $46.7 million[3] via a BEC attack. Yet rising distrust of email, to the point of abandoning email as an everyday business tool, clearly cannot be an option. The consumer and business sectors are too deeply dependent on email.

"BEC spear phishing **cannot be blocked** by legacy email filtering and malware detection technologies"

New Approach Needed

Somehow enterprises must ensure that their employees receive and interact with only authentic and trustworthy messages. An email security solution that establishes per-message authenticity and trust could provide a reliable way to mitigate the risk of targeted email attacks, pushing attacks outside of the circle of trust and reducing criminal effectiveness.

This paper is about just such an alternative approach to securing email – one that can effectively mitigate BEC scams as well as bolster trust in the email ecosystem overall. This can be done by applying machine learning to a large data set in order to get smarter and smarter over time about who is actually sending each email message. The cornerstone for this new approach was set in 2012 with the publication of a new email authentication standard called Domain-based Message Authentication, Reporting & Conformance, or DMARC.

Founding sponsors of DMARC – including Agari, Google, Yahoo, Microsoft, Facebook, JPMorgan Chase and PayPal – sought to collaborate on a new method for combating fraudulent email at Internet-scale. The central idea was to enable email senders to publish easily discoverable email authentication policies, while also enabling receivers to provide authentication reporting to senders.

DMARC has since been embraced by a who's who of top online companies, telecoms, big banks and large enterprises[4] that have been generating mountains of email traffic data using the DMARC specifications. Today DMARC policies protect tens of billions of emails being sent to over 2.5 billion inboxes each and every day. Meanwhile, Agari has applied machine learning analytics to DMARC data sets to improve email authentication infrastructures.

By studying and modeling DMARC data logs, Agari has been steadily gaining a much more cogent understanding about the actual sender of each email. This valuable intelligence infuses Agari Customer Protect™, which is being used by large enterprises to block wide swaths of fraudulent email. On March 29, 2016 Agari introduced a new service, Agari Enterprise Protect™, that can isolate BEC email attacks, in particular, with precise accuracy. Before going deeper on how Agari does this, some context about the infrastructure supporting email is in order.

A Messy Ecosystem

Email traverses an amazingly complex ecosystem that, much like the Internet itself, is an organic hodge-podge of practices and protocols that can be best be described as messy. Email messages hop from one email server to the next; each server along the way naively helps move the message one hop closer to the intended recipient. Accuracy can, at times, be elusive. Authentication of the sender is just not built in.

Phishers continually study this messy arrangement for exploitable vulnerabilities of which there are many. They are adept at tapping into any number of access points where they can insert a spoofed message and send it on its way. And they know full well the sender's authenticity won't be challenged as their message hops from server to server on its way to the recipient.

To defend against this gaping exposure, a multi-billion dollar cottage industry of email filtering and malware detection technologies arose. These two categories of reactive tools are like surveillance cams and x-ray machines. They're terrific at detecting known nefarious email senders, black-listed message types and dangerous malware packages as they arrive at an organization's network gateway, and soon after anything dangerous slips inside the network. They are less effective identifying new iterations of attacks, though they are quick to update blocking lists once fresh intel is obtained.

Spear phishers study and exploit the gaps in this ecosystem. And security vendors watch for novel tactics and continually shore up defenses. From the perspective of an enterprise defending its email system, this translates into a necessity to maintain multiple layers of defenses. A typical enterprise email security portfolio can include several layers of filtering, including placing software in the email gateway to rewrite URLs and sandbox attachments.

A Fortune 1000 enterprise was doing all of this and more in 2014 when the company decided to go a step further and implement the DMARC standard for outbound email and engage Agari to authenticate and protect email they sent on a daily basis to consumers. This was done using Agari Customer Protect™ service.

"As Agari and others bring machine learning to bear on standardized email telemetry data, solutions based on **per-message authenticity and trust** will increasingly keep criminals in check and help to bolster the trustworthiness of the email ecosystem."

Measuring Authenticity

Customer Protect helped the enterprise block millions of instances of fraudulent consumer email originating from unauthorized senders. Building off of that positive outcome, the enterprise in late 2015 began supplying Agari with telemetry data and participating in a project to develop Agari Enterprise Protect™, a new service designed to curtail inbound BEC attacks. The idea behind Enterprise Protect was to apply machine learning analytics to email telemetry data derived from Agari's global email data network.

Agari data scientists built and tested algorithms to parse this standardized email telemetry data. In doing this, they gained a steadily increasing level of clarity about how email moves through its messy ecosystem, hopping from one server to the next. Agari can see, for instance, when a particular sender sends an email from a new domain moments after that domain is initially registered. That's a sure sign of bad guys hustling to stay a step ahead of antimalware crawlers on the hunt -- and seeking to blacklist -- bad domains.

Enterprise Protect identifies and analyzes many other such patterns and can assemble a detailed assessment of all email purporting to be sent from a given enterprise on a given day. That assessment includes assigning an authenticity score to each email message purporting to be from that company.

Meanwhile, Agari's flagship Customer Protect gives organization a clear, actionable daily view of who is sending email in that company's name. For a large U.S. bank, for instance, Agari recently recorded 29,231 email servers sending out email purporting to come from that institution in a 24 hour period. Of those, Agari determined that 332 servers were owned by the bank; 4,641 were operated by legitimate service providers; and 9,732 were run by benign forwarders. The bulk of the remainder – 14,526 email servers – were sending messages not authorized by the bank and thus spoofing their brand.

By applying this type of daily analytics to the Fortune 1000 enterprise's email traffic, Agari helped the company knock down 65 million fraudulent emails in a 12-month period. The result: clients began to open the company's emails more often.

Divining Intent

The challenge for Agari was to take the essence of what machine learning was helping them to see, regarding outbound company's emails, and extrapolate it to inbound email. Agari set out to develop and test statistical estimates of authenticity for email senders not contributing any standardized email telemetry data. This was accomplished by applying what could be learned about how the email ecosystem works in places where Agari did have standardized email telemetry histories.

"Machine learning is learning patterns from a subset of data where you know what the right answers are, and you can recognize patterns, and then translating that to the places where you don't yet know the answers," says Agari's Director of Analytics Scott Kennedy. "It's pattern recognition, basically."

However, broad modeling and assessing of the enterprise's inbound email would only take Agari so far. BEC attacks usually involve one-off messages that can very easily slip through the fire hose of an enterprise's inbound email. So Agari needed to extend its analytics capabilities

"Enterprise Protect assigns an **authenticity score to each email** message purporting to be from that company."

into the arena of checking for valid existing relationships between the sender and the recipient of any given message. For instance, a strong indicator as to whether a sender is who he claims to be turns out to be whether the sender previously has had consistent email exchanges with the recipient.

By driving to this depth of data analytics Agari is moving towards divining intent of the sender. "There are a lot of people out there on the Internet sending messages that look funny for one reason or another, and if you blacklist everything that looks weird you end up with a very high rate of false positives," says Kennedy. "We are trying to pull signals out of that noise that match the characteristics of a true targeted attack, mainly it has to be exploiting identity. It has to be exploiting trust."

The Fortune 1000 enterprise cited in this paper was one of the first organizations to test Agari's new Enterprise Protect service. The company's Chief Security Officer observes that machine learning has helped his organization keep pace with spear phishers as they change tactics. "It's a set of controls with constantly improving filtering capability based on feedback that we provide in terms of what works and what doesn't," he says. "When we do find a targeted phishing email that does get through, then we know the bad guys are changing their tactics and we can add different heuristic information into a revised algorithm."

On one particular day last fall, Enterprise Protect identified and alerted this enterprise to five BEC emails sent to its CEO, solid confirmation of the efficacy of Agari's algorithms, the CSO says. And there was an unexpected added benefit: the company's security staff gained valuable counter intelligence about the parties attacking the company.

"We analyze every single email we block primarily because we want to know which threat actor or what category of threat actor is originating in the emails," the CSO says. "That's useful just from an intelligence standpoint because we can follow the migration of certain threat actors as the move from one channel to another channel."

This enterprise's experience with Customer Protect and Enterprise Protect demonstrates the upside of taking a new approach to email security that revolves around ensuring that customers and employees receive and interact with only authentic and trustworthy messages. Wide adoption of standardized email telemetry, and Agari's global email data network provide a solid foundation for accelerating this approach. As Agari and others bring machine learning to bear on standardized email telemetry data, solutions based on per-message authenticity and trust will increasingly keep criminals in check and help to bolster the trustworthiness of the email ecosystem.

"This new approach to email security revolves around ensuring that customers and employees receive and interact with only **authentic and trustworthy messages.**"

1. September 2014 Panel Survey, Pew Research Center
2. http://www.ic3.gov/media/2015/150827-1.aspx
3. http://www.itgovernanceusa.com/blog/ubiquiti-fraud-the-46-million-cyber-crime/
4. https://dmarc.org/about/history/