Return Path

Email Threat Intelligence Report

Insights Into Brand Spoofing Tactics

September 2015

Authors:

Rob Holmes
Matt Moorehead
Ben Cassedy
Liz Dennison

returnpath.com/stopemailfraud

Table of Contents

Introduction 3

Three Key Hypotheses 4

1: Fraudsters use snowshoe spamming in large phishing attacks 4

2: Rotating elements of subject lines to appear personalized 4

3: Spoofing the Display Name 4

Methodology 5

Data Analysis & Findings 6

1: There is no discernible pattern to snowshoe spamming relative to the size of the attack 6

2: Fraudsters do not rotate elements of their subject lines 7

3: The most frequently spoofed Header From field is the Display Name 9

Conclusion 12

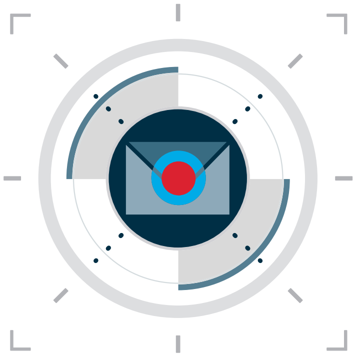

Introduction

Email fraud is on the rise—up more than 162% from 2010-2014. To defend customers, brand reputation, and revenue, organizations around the world are implementing the authentication standard DMARC (Domain-based Message Authentication Reporting and Conformance) to block bad email before it reaches consumer inboxes.

But as email authentication protocols get more advanced, so too do fraudsters.

There's no doubt that implementing DMARC is a great first step. But it's not a complete solution, protecting your brand from only 30% of email threats, according to research previously conducted by Return Path.

Defending against the other 70% requires a comprehensive understanding of the tactics fraudsters use to bypass email authentication.

To gain that understanding, we tapped into the Return Path Data Cloud, which empowers us to detect, classify, and analyze 6 billion emails every day from more than 70 mailbox providers.

Our main objectives were to:

1. **Better assess threats beyond DMARC**
2. **Better understand the tactics used by fraudsters to bypass DMARC**
3. **Identify ways to combat those tactics directly**

Display Name	**The name that will be displayed on the "From:" line in emails sent or received.**
Header From	**The address contained in the From: field of an email, which is visible to all email users (aka "friendly from).**
Email Domain	**The domain pulled from the header of each message. The From domain includes the email address the message was sent from. The Email Domain is the portion of the From address right of the @ symbol.**
Email Name	**Element to the left of the @ on the Header From email.**
Subject	**Subject line field in an email.**

Three Key Hypotheses

The email threat landscape is a constantly evolving one. Fraudsters implement a variety of tactics to launch phishing campaigns—too many to highlight in a single report. For this analysis, therefore, we chose to focus on three key tactics we suspect fraudsters use to circumvent email authentication mechanisms like DMARC, and test those suspicions against empirical threat data.

1 Hypothesis 1: Fraudsters Use Snowshoe Spamming in Large Phishing Attacks

One tactic we suspect fraudsters use is snowshoe—or distributed—attacks.

Just as a snowshoe spreads the load of a person's weight across a wide area of snow, snowshoe spamming distributes spam from various IP addresses in order to dilute reputation metrics, evade filters, and avoid getting blacklisted.

Traditional spam filters struggle with snowshoeing because they may not see enough volume from a single IP to trigger the filter. Therefore, we suspect fraudsters use this technique in large-scale phishing attacks to stay under the radar of volume-based and IP-based filtering.

2 Hypothesis 2: Fraudsters Rotate Elements of Subject Lines to Appear Personalized

Marketers know that personalizing email is the best way to boost email engagement.

According to Experian, personalized emails improve transaction rates by 600 percent. Specifically, personalized subject lines deliver 26 percent higher unique open rates overall, with travel companies experiencing the "biggest boost."

We suspect fraudsters know this too—and mimic elements of the personalization that works so well for marketers by serializing subject lines.

They may include a fake shipping order or tracking number in a subject line, changing that number just enough so mailbox providers can't detect that the same fraudulent email is going to thousands of victims.

Or, if a fraudster is phishing a bank, they may include a fake account number in a subject line, using the same rotating technique to coast under the radar of email filters.

3 Hypothesis 3: Fraudsters Spoof the Display Name

Spoofing the Display Name of a legitimate brand is, we suspect, one of the easiest ways for fraudsters to appear legitimate and bypass email authentication protocols like DMARC by using a sending domain that is not owned by the brand targeted in the attack.

For example, if a fraudster were spoofing the hypothetical brand "My Bank," the email may look something like:

To: **You** <you@yourdomain.com>
From: **My Bank** <accounts@secure.com> ←
Subject: **Unauthorized login attempt**

Since My Bank doesn't own the domain "secure.com," DMARC won't block this email on My Bank's behalf, even if My Bank has set their DMARC policy for mybank.com to reject messages that fail to authenticate.

This fraudulent email, once delivered, may appear legitimate because most user inboxes only present the Display Name.

Methodology

For this project, we leveraged the Return Path Data Cloud—our proprietary network of over 70 mailbox and security providers representing 2.5 billion email accounts and in-depth behavioral insights from more than 2 million individual consumer inboxes.

During a 40 day period (July and August 2015), we analyzed over 240 billion emails from more than 100 data feeds to identify what threats were leveraging brands and targeting their customers. We also queried Return Path Sender Score to detect bad IP addresses sending email on behalf of legitimate brands—no matter how low the volume.

To focus on the specific tactics fraudsters used, we selected 40 of the world's largest brands across various industries that, historically, have been subject to email fraud.

Looking at the email threat ecosystem, we then applied filters to uncover attacks pretending to come from the 40 identified brands and processed those results through Return Path's proprietary threat classification algorithm to eliminate false positives. This revealed 760,000 malicious emails, targeting the brands, which were then used for the analysis of this report.

Industry	Brands
Banking	12
Healthcare/Insurance	5
ISP/Telco/Technology	8
Retail/eCommerce/Gaming	9
Social Media	3
Other	3
Total Brands	40

Data Analysis & Findings

1 Finding 1: There Is No Discernible Pattern to Snowshoe Spamming Relative to the Size of the Attack

To assess the validity of our hypothesis that snowshoe spamming is a common technique used by fraudsters in large phishing attacks was correct, we collapsed the 769,792 sample email threats into 51,361 unique fraudulent campaigns. A "campaign" is defined as a unique combination of a brand being targeted with a given phishing or malware template.

Within this data set, there was a very long tail of small attacks—80% of fraudulent emails analyzed could be attributed to just 4.2% of the campaigns identified. We excluded these smaller attacks from our subsequent analysis, and focused on the top 100 campaigns that accounted for 33.4% of the fraudulent emails we identified targeting our 40 brands.

We classified these attacks as "Huge," "Large," or "Medium" using the following intervals relative to our email sample size:

Volume of sample fraudulent emails seen	Attack size	Number of attacks observed
>7,500	Huge	9
>2,500	Large	15
>500	Medium	76
Total		**100**

We then looked at how attacks were distributed across IP addresses. If more than 90% of messages observed came from a single IP address, the attack was classified as "not distributed." We applied additional intervals to our sample data based on the following criteria:

IP addresses observed in the attack	Level of attack distribution	Number of attacks observed
>300	Very High	22
>100	High	12
>20	Medium	14
<20	Low	25
Not distributed	Not distributed	27
Total		**100**

Next, to uncover whether or not there is a correlation between the size of an attack and the nature of its distribution, we assessed these two variables against one another:

	Attack size			
Attack distribution	Huge	Large	Medium	**Total**
Very High	0	0	22	**22**
High	2	1	9	**12**
Medium	3	6	5	**14**
Low	0	3	22	**25**
Not distributed	4	5	18	**27**
Total	**9**	**15**	**76**	**100**

We found no discernible pattern that links the size of an attack to how distributed it is. There may be other factors that play a role in predicting the fraudsters' preferred attack methods (for example, the vertical or brand targeted).

Data Analysis & Findings (continued)

Our key conclusions from this inquiry into snowshoe spamming were:

1. Law enforcement has made significant inroads in shutting down big botnets over the past year. However, there is evidence that fraudsters still use botnets to launch not just spam but also distributed phishing or snowshoe attacks—34 of the 100 largest attacks were launched from a highly distributed network of IP addresses.

2. There is no discernible pattern in our current sample to confidently predict that the biggest phishing attacks are launched on a distributed IP address basis: 4 out of the 9 biggest attacks we observed were not distributed, while 22 of the 76 medium-sized attacks were sent from a very distributed set of IP addresses (more than 300 in our sample data set alone).

3. Assessing IP reputations should therefore continue to provide value in defending against some large-scale phishing attacks.

2 Finding 2: Fraudsters Do Not Rotate Elements of Their Subject Lines

To assess the validity of our hypothesis that fraudsters in some way mimicked elements of personalization of subject lines to avoid detection, we analyzed the subject lines of the 769,792 malicious email messages.

We found that the overwhelming majority of fraudulent emails—98.15%—*did not* have personalized subject lines. Most fraudsters simply do not go to the trouble of serializing subject lines, likely because they don't want to expend time or cost unnecessarily.

The most popular subject lines *did* incorporate a sense of urgency to compel recipients to open the email.

The top three subject lines were:

Subject Line	% of Email Messages
Account Verification	12.06%
Account Notification	7.49%
Your Email Account Has Been Suspended	6.18%

Data Analysis & Findings (continued)

Of the 1.85% of the email messages that *did* have quasi-personalized subject lines, we found they fell under four interesting themes:

1. Social Media Scams

Fraudsters personalized subject lines in social media scams by rotating the number of alerts, e.g., *"You Have 6 New Friend Requests."* These social media alerts are a great hook for fraudsters not only because they are inherently personalized, but also because it's easy to serialize the alert numbers to help avoid detection.

2. Account Security

Fraudsters also personalized subject lines pertaining to "account security" emails. These subject lines suggested that unauthorized users were signing in to the recipients' legitimate accounts. Fraudsters serialized the email addresses in subject lines such as: *"New sign-in from (random email address) account security!"*

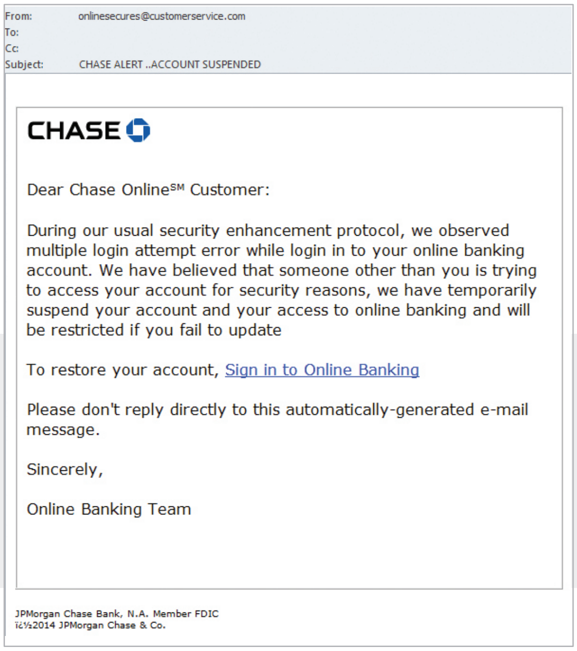

From: onlinesecures@customerservice.com
To:
Cc:
Subject: CHASE ALERT ..ACCOUNT SUSPENDED

CHASE

Dear Chase Online℠ Customer:

During our usual security enhancement protocol, we observed multiple login attempt error while login in to your online banking account. We have believed that someone other than you is trying to access your account for security reasons, we have temporarily suspend your account and your access to online banking and will be restricted if you fail to update

To restore your account, Sign in to Online Banking

Please don't reply directly to this automatically-generated e-mail message.

Sincerely,

Online Banking Team

JPMorgan Chase Bank, N.A. Member FDIC
ï¿½2014 JPMorgan Chase & Co.

3. Calls to Action with Reference Number

The third tactic fraudsters implemented was to combine calls to action with random reference numbers or timestamps. For example: *"Paypal : Your Account will be Limited, Please Update Your Account ! 15/08/2015 06:16:06".*

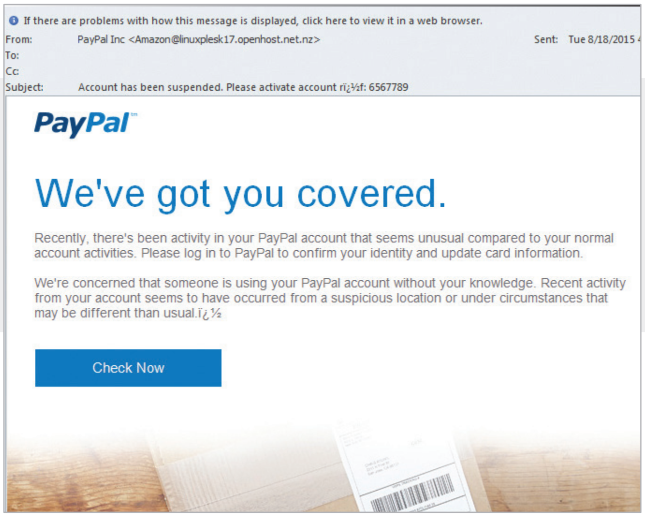

If there are problems with how this message is displayed, click here to view it in a web browser.
From: PayPal Inc <Amazon@linuxplesk17.openhost.net.nz>
Sent: Tue 8/18/2015
To:
Cc:
Subject: Account has been suspended. Please activate account rï¿½f: 6567789

PayPal

We've got you covered.

Recently, there's been activity in your PayPal account that seems unusual compared to your normal account activities. Please log in to PayPal to confirm your identity and update card information.

We're concerned that someone is using your PayPal account without your knowledge. Recent activity from your account seems to have occurred from a suspicious location or under circumstances that may be different than usual.ï¿½

4. HR Scams

The final theme we discovered were HR alerts. Fraudsters used subject lines that coupled job offers with serialized reference numbers, e.g., *"Job offer (A123B456C)".*

While these subject line personalization themes are interesting to note, the data we uncovered decidedly disproves our initial hypothesis that subject line serialization was a popular tactic to bypass filters. On the contrary, fraudsters prefer a more template-based approach.

Data Analysis & Findings (continued)

3 Finding 3: The Most Frequently Spoofed Header From Field Is the Display Name

To test our Display Name spoofing hypothesis, we wanted to know not only if fraudsters were spoofing Display Names (we were pretty sure of that) but also how they were doing it.

Of the 769,792 malicious email messages we identified, 44.23% featured the brand in the Display Name—14.30% were an exact match of the brand, and 29.93% were a wildcard match, e.g., the brand name was contained within the full Display Name.

Display Name/Brand relationship	% Threats
Exact match	14.30%
Wildcard match	29.93%
Other	55.76%
Total	100.00%

As we suspected, nearly half of all email threats spoofed the brand in the Display Name.

But the Display Name is only one element of the Header From: the other is the Header From email address. Looking at the email address, we analyzed both the Email Name (to the left of the @) and the Email Domain (to the right of the @) and discovered that nearly 30% of threats spoofed the brand in the email address. Of those threats, more than two thirds focused on spoofing of the Email Domain alone:

Email Domain spoofing analysis:

		Email Domain		
		Brand spoofed	Brand not spoofed	Total
Email Name	Brand spoofed	0.98%	7.52%	8.51%
	Brand not spoofed	20.57%	70.92%	91.49%
	Total	21.56%	78.44%	100.00%

Data Analysis & Findings (continued)

When we looked at the union of Display Names and email addresses, we discovered the following spoofing behaviors in relation to the Header From field:

Header From spoofing analysis:

		Email Address		
		Brand spoofed	Brand not spoofed	Total
Display Name	Brand spoofed	10.63%	33.61%	44.24%
	Brand not spoofed	18.46%	37.31%	55.77%
	Total	29.09%	70.92%	100.00%

In the majority (62.69%) of email threats, fraudsters spoof elements of the Header From field, the most popular being the Display Name field (for which there is currently no authentication). So, what of the remaining 37.31% of threats for which there were no clues in the Header From field?

We analyzed the subject field of those emails for which there was not spoofing of the Header From and determined that, in the majority of cases, the subject field did not offer clues as to the brand that was being targeted.

Subject included brand	% Threats
Yes	7.45%
No	92.55%
Total	100.00%

Our assessment of Header From spoofing against Subject spoofing revealed that headers give a strong indication of spoofing (65.47% of messages featured spoofing in the headers), with the Header From giving the clearest indicator of threat (58.78% of from headers were spoofed). This isn't surprising, since identity deception is a key component of email fraud and the sending identity is the Header From field.

Header spoofing analysis:

		Subject Header		
		Brand spoofed	Brand not spoofed	Total
Header From	Brand spoofed	3.91%	58.78%	62.69%
	Brand not spoofed	2.78%	34.53%	37.31%
	Total	6.69%	93.31%	100.00%

Data Analysis & Findings (continued)

Only in 34.53% of email threats analyzed were those clues buried exclusively in the body of the emails. On closer inspection of the 265,817 (34.53%) emails that only featured the brand in the body of the email, even a cursory look at the headers suggests an overwhelming amount of advanced fee fraud:

1. Looking at the From headers, James B. Comey alone seems to have sent a whopping 7.60% of those fraudulent messages (James B. Comey is the Director of the FBI) using email addresses ranging from "James B. Comey, Jr." <testing@gpu-rb.ru> to "JAMES B. COMEY" <FBI-ALERT@FBI.GOV>. (If only everyone implemented DMARC!)

2. Inspecting the Subject headers, "Good News !!!" and "ATTENTION BENEFICIARY" clearly point to advanced fee fraud, while the popular subject of "Federal Bureau of Investigation (FBI)" suggests rather more ominous scams.

From this analysis, five key conclusions emerged:

1. Brand identities are regularly spoofed to get into the inbox and dupe people into opening fraudulent email.

2. The Header From field is most often used to spoof brand identities and the most frequently spoofed element is the Display Name, for which there is currently no authentication mechanism.

3. The Subject header is also sometimes used to spoof identities, but rarely in isolation.

4. Domain owners should nevertheless implement DMARC to make it easier for receivers of email to differentiate between good and bad email.

5. With phishers making use of a variety of angles to target victims, brands need to educate their customers about the risks associated with not only phishing and spoofing attacks, but also all types of email fraud.

Conclusion

For this analysis, our goal was not to uncover each and every tactic fraudsters use to launch phishing attacks—such an endeavor would be futile in the ever-shifting email threat landscape we've come to know so well. We did, however, seek to gain a deeper understanding of that landscape by testing three key assumptions against the email threat intelligence at our disposal.

After analyzing more than 760,000 email threats, we found that:

While there is no discernible pattern to snowshoe spamming, this method is still rife and monitoring IP reputations needs to be part of a multi-faceted email fraud protection strategy.

Fraudsters do not go to the trouble of rotating elements of their subject lines, preferring a more template-based approach. Access to message-level data from email threat intelligence sources should help you prioritize your efforts around attack mitigation.

The most frequently spoofed Header From field is the Display Name, for which there is currently no authentication mechanism. Visibility into Display Name spoofing is critical in identifying and responding to phishing attacks leveraging your brand.

The learnings from this study can inform email fraud protection at your company in two ways.

First, prioritize DMARC implementation—it's the most direct way to keep bad email out (and the good email in) of consumer inboxes.

Second, the more you know about the nature of email attacks spoofing your brand, the better. As our analysis proves, fraudsters like to mix and match tactics to reach their victims. While DMARC is a great first step, it is not enough. Protect your brand from the 70% of email threats beyond DMARC by studying their anatomy. Only then can you implement the right suite of solutions to fight back.

Return Path

Getting Started With DMARC

Get Your Copy at:
returnpath.com/research/getting-started-with-dmarc

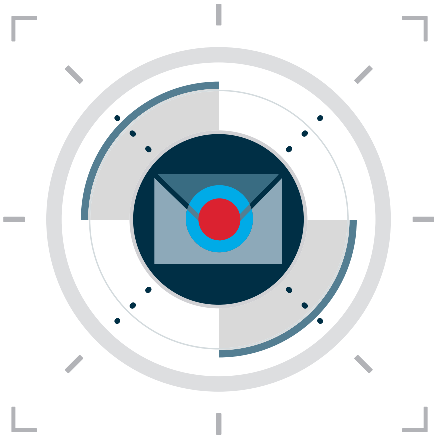

About Return Path

Return Path analyzes the world's largest collection of email data to show businesses how to stay connected to their audiences, strengthen their customer engagement, and protect their brands from fraud. Our data solutions help analysts understand consumer behavior and market trends. We help mailbox providers and security providers around the world deliver great user experiences and build trust in email by ensuring that wanted messages reach the inbox while spam and abuse don't.

Find out more about Return Path Email Fraud Protection at
www.returnpath.com/stopemailfraud
@stopemailfraud

Return Path