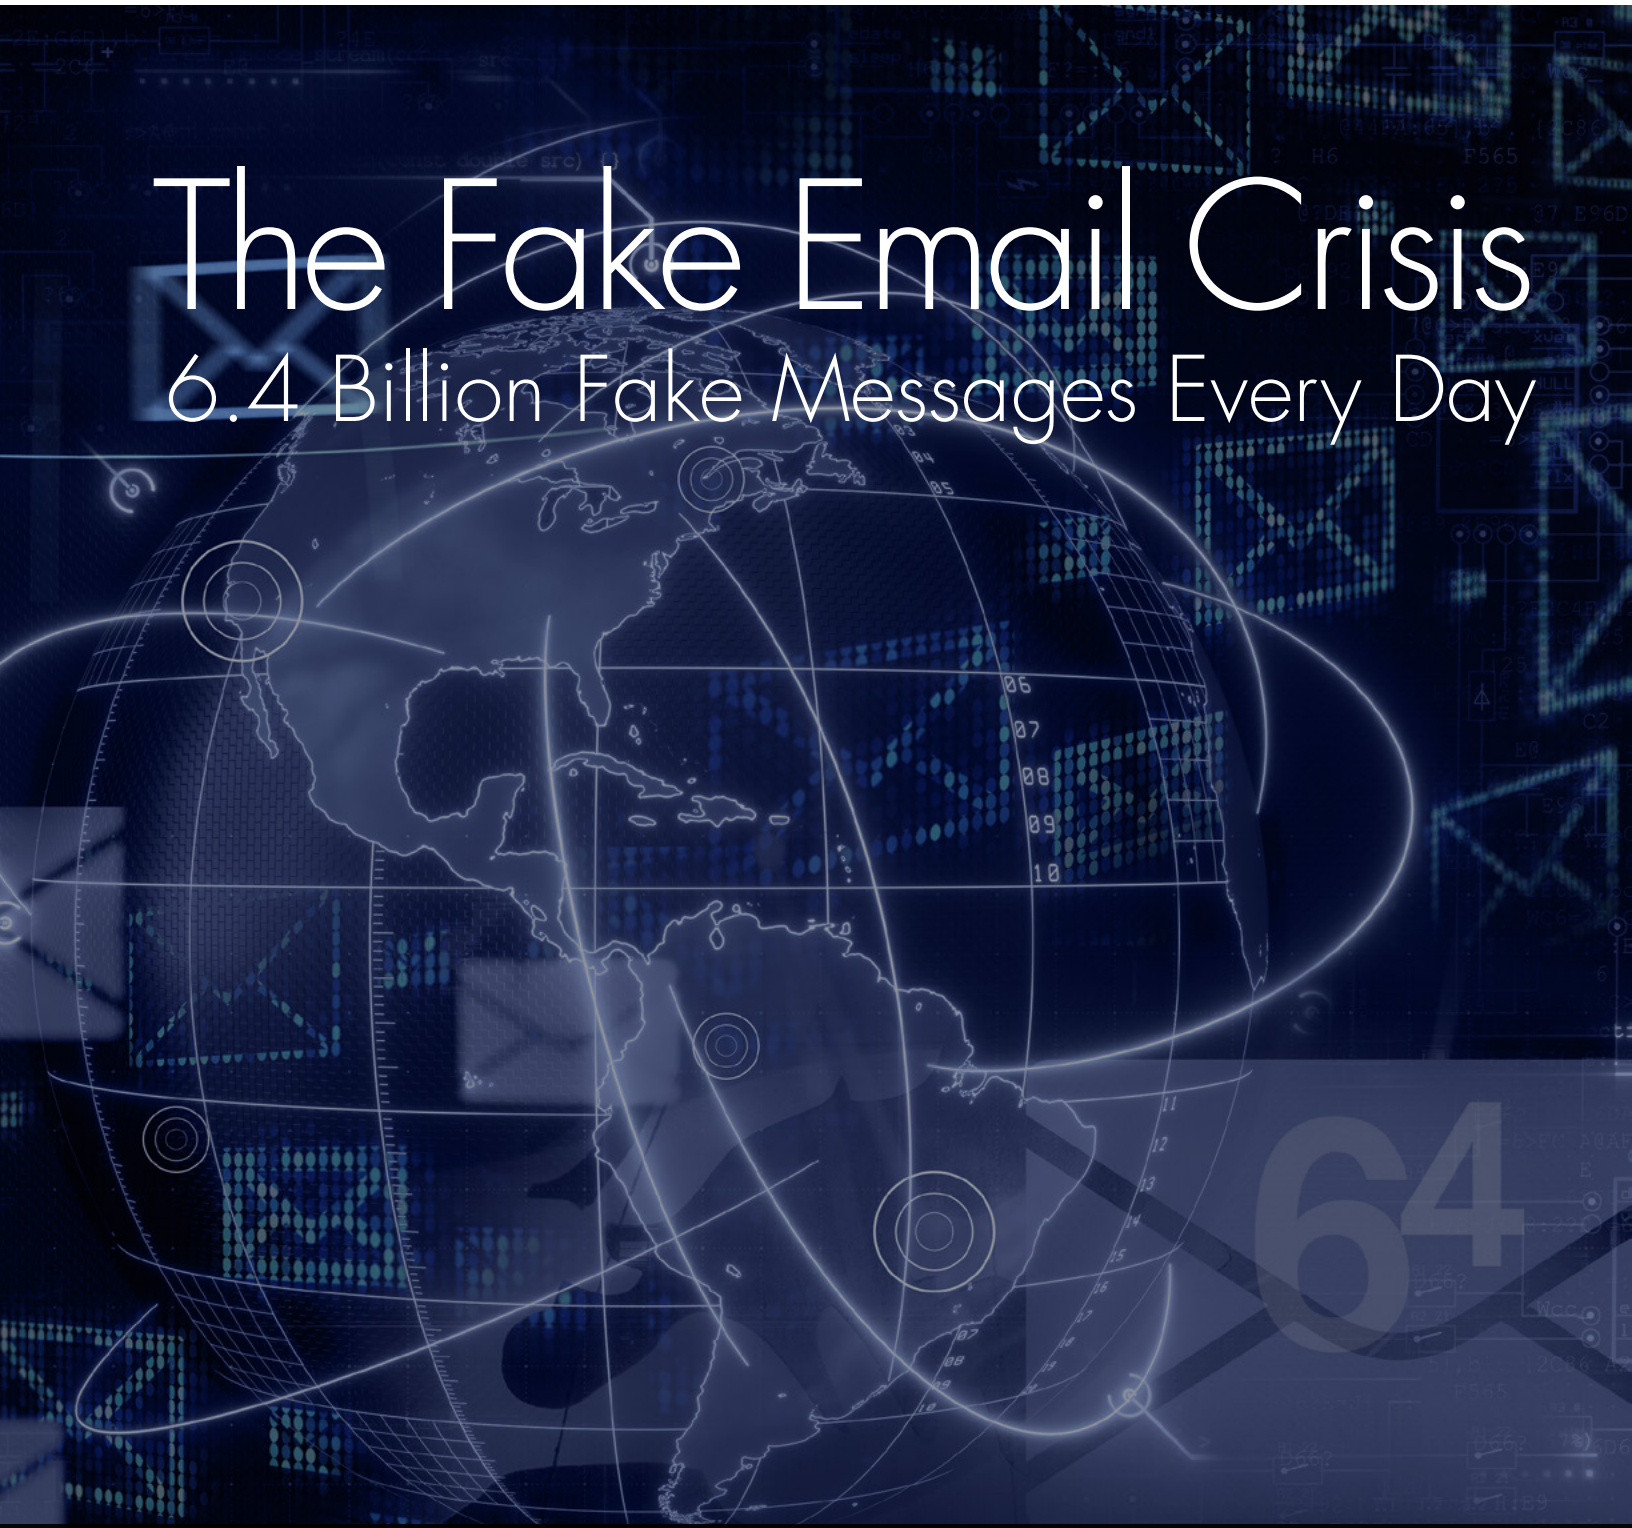

# The Fake Email Crisis

## 6.4 Billion Fake Messages Every Day

## Email Fraud Landscape, Q2 2018

# Executive Summary

The crisis of fake email continues. Far from being merely a "social engineering" issue, fake email is a direct result of technical issues with the way email is implemented: It lacks a built-in authentication mechanism making it all too easy to spoof senders. However, this problem is also amenable to a technical solution, starting with the email authentication standards DMARC, SPF, and DKIM.

For the purposes of this report, Valimail used proprietary data from our analysis of billions of email message authentication requests, plus our analysis of more than 3 million publicly accessible DMARC and SPF records, to compile a unique view of the email fraud landscape. Now in its third consecutive quarter, our report shows how the fight against fake email is progressing worldwide, in a variety of industry categories.

## Key Findings

- 6.4 billion fake emails (with fake From: addresses) are sent worldwide every day
- The United States continues to lead the world as a source of fake email
- The rate of DMARC implementation continues to grow in every industry
- DMARC enforcement remains a major challenge, with a failure rate of 75-80 percent in every industry
- The rate of SPF usage continues to grow in every industry
- SPF errors remain a significant problem
- The U.S. federal government leads all other sectors in DMARC usage and DMARC enforcement

## Life on Planet Email

Email continues to be a robust, effective medium for communications worldwide, and it is both the last remaining truly open network in wide use as well as the largest digital network, connecting half the planet.

Half the global population has at least one email account, making email's reach larger than any other digital network by far. The Radicati Group estimates that 3.8 billion people worldwide use email, with an average of 1.75 accounts per user, for a total of 6.7 billion email accounts in use.[1]

Those accounts are not merely sitting dormant, either. Around the world, Radicati estimates that 281 billion email messages are sent and received every day.

**7.6 Billion** - World Population in 2018
**3.8 Billion** - Number of Email Users

Source: Radicati Group

Why has email continued to defy predictions of its death? Simply put, email works. As a truly open network, it is based on open protocols and a completely distributed architecture, meaning no single company or country "owns" email. It provides a common medium for the world to communicate, regardless of location, OS, connectivity, or language. What's more, businesses have found that email is incredibly effective at communicating with customers, both for B2B and B2C marketing, and studies abound showing the high ROI of using email.[2]

## The Fake Email Crisis

**91%**

OF ALL CYBERATTACKS START WITH A PHISHING EMAIL

Source: PhishMe

Yet email is under attack. With no authentication built in to email's core protocols, it has always been trivial for hackers to impersonate others: Your boss, your CFO, a trusted business partner, even government agencies.

In the early days of email, prior to the 1990s, impersonation was not a huge problem (though it was used for the occasional prank by academic computer scientists).

Once the Internet entered its commercial era in the mid-1990s, however, impersonation started to become more prevalent as a technique used by spammers to obscure their identities. But since the early 2000s, spam filters have largely solved that problem, relegating most such messages to users' spam folders, where they can be safely ignored.

However, impersonation has moved to a new field: Fake emails as a vector for system compromise. Multiple studies have identified spear phishing as a primary attack technique used in the vast majority of security incidents, including account compromise, exfiltration, penetration, and more: It plays a role in 90 percent or more of all cyber attacks.[3] The cost of just one type of fake-email attack, the business email compromise (or BEC), has exceeded $12 billion since 2013, according to the FBI.[4]

Now, it is true that not every phishing mail is an impersonation: Sometimes the apparent sender is unknown to the recipient, so no impersonation is necessary (e.g. in cases where the sender poses as a job seeker and attaches a malware-laden resume). However, in the vast majority of spear phishing cases some kind of identity deception is in play, with the sender using a fake From: address, a deceptive "Friendly from" name, or a deceptive "lookalike" domain. We refer to all of these as impersonation attacks or fake emails.

In short, impersonation — fake email — is a core hacking technique. Relegating it to the category of "social engineering" and giving it the virtual equivalent of a shrug is not a sufficient response. There are better ways.

1. *Radicati Email Statistics Report*, 2018-2022
2. *Email Continues to Deliver Strong ROI and Value for Marketers* (eMarketer, September 2016); *70 Email Marketing Stats Every Marketer Should Know* (Campaign Monitor).
3. *Proofpoint*, April 2018; *IronScales 2017 Email Security Report; PhishMe*, December 2016
4. *Business Email Compromise Losses Top $12B, FBI Warns* (eWeek, July 2018)

# Email Fraud Frequency

Data on the frequency of fake emails can be hard to come by. However, an analysis of a representative subset of the message authorization requests processed by Valimail provides a snapshot of one type of impersonation: The fake From: address.

Valimail Enforce, our email authentication automation solution, provides real-time responses to mail gateway requests for Domain-based Message Authentication, Reporting & Conformance (DMARC), Sender Policy Framework (SPF), and DomainKeys Identified Mail (DKIM) requests. We also collect and analyze aggregate reports generated by these mail gateways. Valimail processes many billions of messages on behalf of our customers, and as a result, we have a unique view of the fake email universe.

## MESSAGE DISPOSITION IN **H1 2018**

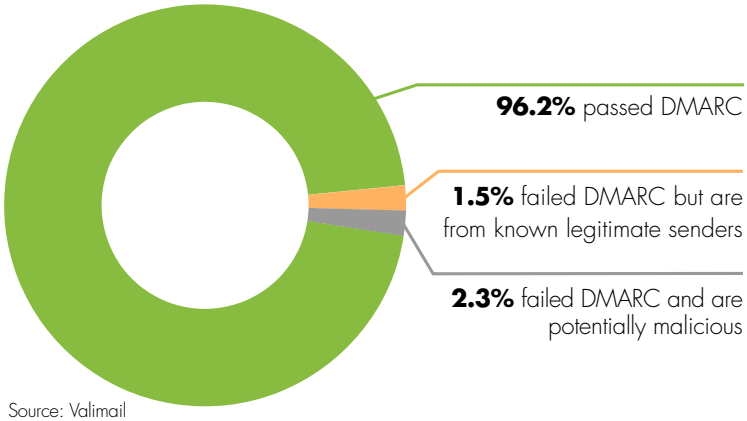


Source: Valimail

In the first half of 2018, about 96 percent of the messages processed by mail gateways on behalf of our customers passed DMARC (they were identified as legitimate users of their apparent sender's domain name). Another 1.6 percent failed DMARC, meaning their senders were not authorized by the domain owner but were from senders known to be legitimate. These DMARC failures are attributable to new customers who have not yet authorized all the cloud-based services that should be able to send email on their behalf.

The remaining 2.2 percent of messages in H1 2018 failed DMARC and come from senders we categorize as suspicious or "possibly malicious."

## This may not sound like a large percentage, but:

1. considering the enormous volume of messages, it would mean that **6.4 billion fake emails are being sent every day worldwide**.
2. Valimail's dataset is dominated by customers who have implemented DMARC at enforcement for months or years. We have observed that as customers implement DMARC enforcement, the number of exact-domain impersonation attempts tends to fall off as attackers realize their messages are no longer being delivered.
3. this total considers only exact-domain spoofing (fake From: addresses), not friendly-from or lookalike-domain attacks. However, exact-domain fakes are the hardest to detect and can cause the most lasting damage.

We have noticed that fraud rates go up and down as attackers around the world launch new phishing campaigns or discontinue them. In our Email Fraud Landscape for Q1 2018, we reported that about 5 percent of total message volume in 2017 was suspicious and a whopping 1 in 5 messages in the month of October were fraudulent. The lower rate in the current report is not necessarily a sign of continuing progress but rather a temporary anomaly. We will learn more in future quarters as this number goes up or down.

**In short, however, fake email is a significant problem.**

## A World of Fake Emails

Looking at the sources of fake email worldwide, Valimail has found a similar pattern to what we observed last quarter: The leading source remains the United States with Canada and the Russian Federation remaining prominent sources of fake email. The United Kingdom and Vietnam now appear in the top 10, while Canada, France, and Germany have moved further down the list. Thailand and the Netherlands have now dropped out of the top 10.

## Top 10 Sources of Suspicious Email, Q2 2018

Source: Valimail

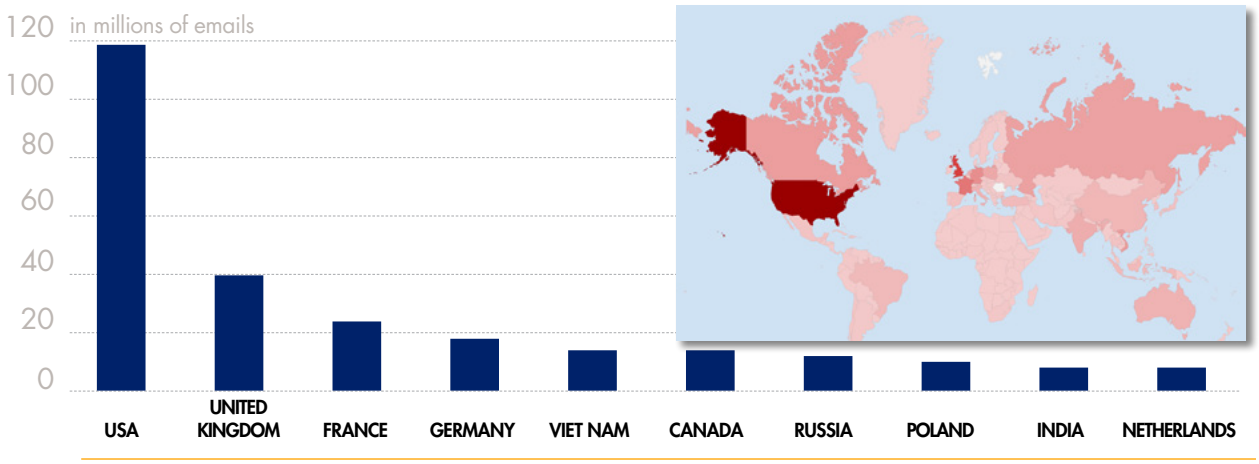


## The Solution to Exact-Domain Fakes: DMARC

There is an answer to the fake email crisis: Domain-based authentication. A cornerstone of that solution is already in place: It's the DMARC standard mentioned above.

DMARC, when properly implemented and configured to an enforcement policy, provides complete authentication and fraud protection against exact-domain fakes. DMARC:

- instructs receiving mail servers to perform certain authentication checks: Namely, SPF and/or DKIM authentication.
- requires that the address shown in the From: field of a message match the domains used in SPF or DKIM authentication. This is known as "alignment" and is the key to making DMARC an anti-fraud tool — not just an authentication or deliverability component.
- provides domain owners with the ability to specify what should happen when a message fails authentication (do nothing and deliver as normally, send it to a spam folder, or delete it entirely).
- requires mail gateways to report the results, usually in aggregate, to the domain owner or a designated agent.

In this scheme, it takes two to tango: Domain owners must publish DMARC records if they want to be protected, and mail gateways/inbox providers must check for the presence of DMARC records, and do authentication checks, for all incoming messages.

The good news is that virtually all major inbox providers worldwide now do DMARC checks on all incoming messages. Valimail's analysis shows that 75 percent of inboxes worldwide (5 billion in all) will do DMARC checks and enforce the domain owners' policies if those policies exist. That number has not significantly changed since last quarter.

## DMARC Support Among Inbox Providers Worldwide

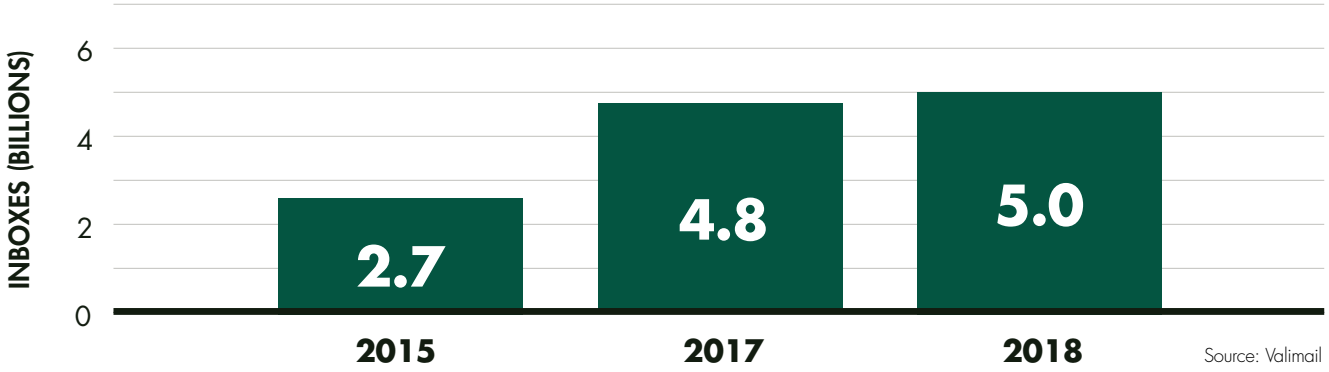


Source: Valimail

Recently, the number of domain owners publishing DMARC records has dramatically increased. The latest figures available from Farsight and DMARC.org show a threefold increase in the total number of domains with DMARC over the course of 2017.[5]

## Total DMARC Records and % Change by Month

Source: DMARC.org

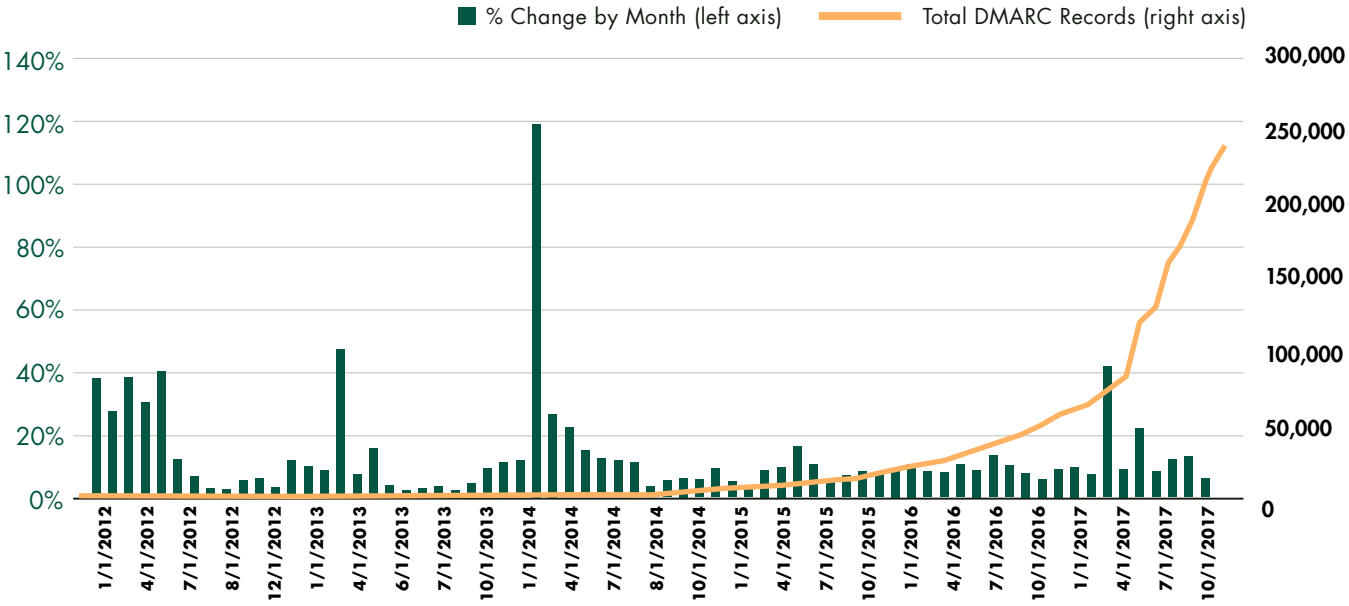


Already, many inbox providers (such as Gmail and Microsoft) will flag messages that fail authentication or which lack authentication with question marks or red warning signs. Refusing delivery for non-authenticated messages is the next step. Some experts in the email world are talking about a global shift in the near future to a "no auth, no entry" policy, in which email is not delivered unless it passes DMARC authentication.

In the rest of this report, we will look at the DMARC adoption rates among various global industries.

5. *Farsight DNS Data* (DMARC.org, January 2018).

## Email Authentication Continues to Grow

Valimail regularly queries over 3 million domains for the presence of published DMARC and SPF records, and performs detailed analysis on any records that we find. (It is not possible to do DKIM analysis in this way, since DKIM records are only accessible through specific "selectors" specified in individual email messages.)

For this report, we examined the DMARC records published by thousands of organizations worldwide, grouped into 11 different categories. In most of these cases we only tested the company's primary domain. For most of these categories we now have data from three successive quarters, which provides a revealing window not only on how these industries compare to one another, but also how they are changing over time.[6]

## DMARC Usage

Source: Valimail

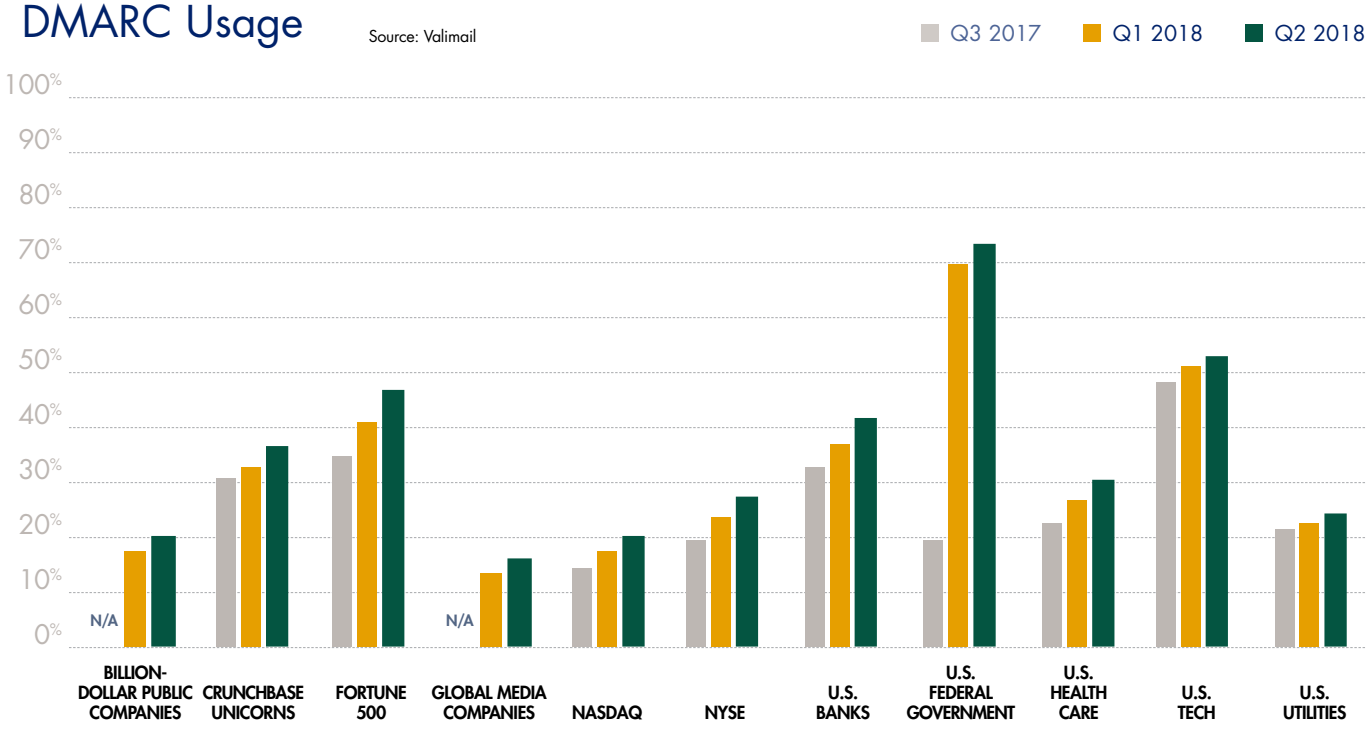


In virtually every category Valimail has examined, we've seen a steady increase in the number of domains that have published DMARC records of any kind.

There is one exception: In the cohort of U.S. federal government domains, we've seen a dramatic, unprecedented increase. This is due directly to the Department of Homeland Security's October, 2017 directive requiring all executive-branch agencies to implement DMARC on a one-year timeline. Since the executive branch accounts for the vast majority of the 1,315 federal .gov domains, that directive, known as BOD 18-01, has had a huge impact on DMARC usage in this group. In Q3, just as BOD 18-01 was issued, 18.5 percent of federal domains had DMARC records; the total is now a massive 71.1 percent.

Two other sectors — U.S. tech companies with $1 billion or more in revenues, and the Fortune 500 — now show more than 40 percent adoption in their use of DMARC. A third category, U.S. banks with $1 billion or more in revenues, is just touching 40 percent.

6. For this report, we updated the list of domains for the Crunchbase unicorns and Fortune 500 categories, to ensure that we were using the most current list of companies for each. The other categories' lists remain the same as in our previous two reports.

## Enforcement Failure Rate

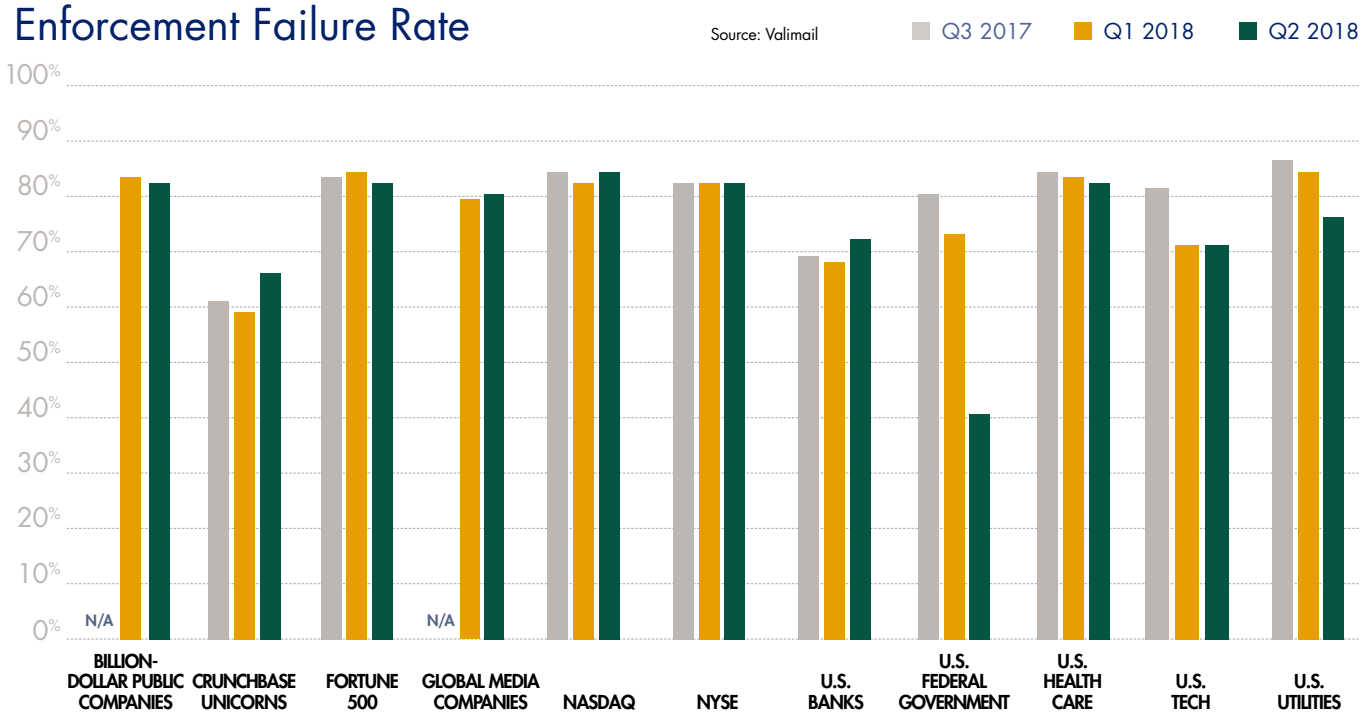


Publishing a DMARC record is one thing, but configuring it correctly and completely is another. Domain owners must ensure that all cloud-based services that send email are duly authorized. They need to ensure that the DMARC and SPF records are all correctly configured and then switch their DMARC policy to enforcement if they wish to realize the standard's anti-impersonation benefits.

To date, most companies that attempt DMARC do not complete the journey. The enforcement failure rate — the percentage of companies that deploy a DMARC record but don't get to enforcement — hovers around 75-80 percent for almost every category of company we have studied. While that number has decreased slightly over the past few quarters in a few categories (reflecting incremental improvements at getting to enforcement), the failure rate has remained fairly stable over the past three quarters.

Two categories stand out for having slightly lower failure rates (and thus, higher DMARC enforcement success): Crunchbase unicorns show a failure rate of only 65 percent this quarter, and the U.S. federal government has a failure rate around 40 percent — the lowest (i.e. best) we've ever seen.

Again, the success of the U.S. government is attributable to compliance pressure: BOD 18-01 requires agencies to get their DMARC policies to enforcement (specifically, a "reject" policy) by October 16, 2018, and many agencies have responded by doing exactly that — particularly with domains that aren't actively used to send email, and which are easier to configure as a result.

## Fraud Protection (Enforcement)

Source: Valimail

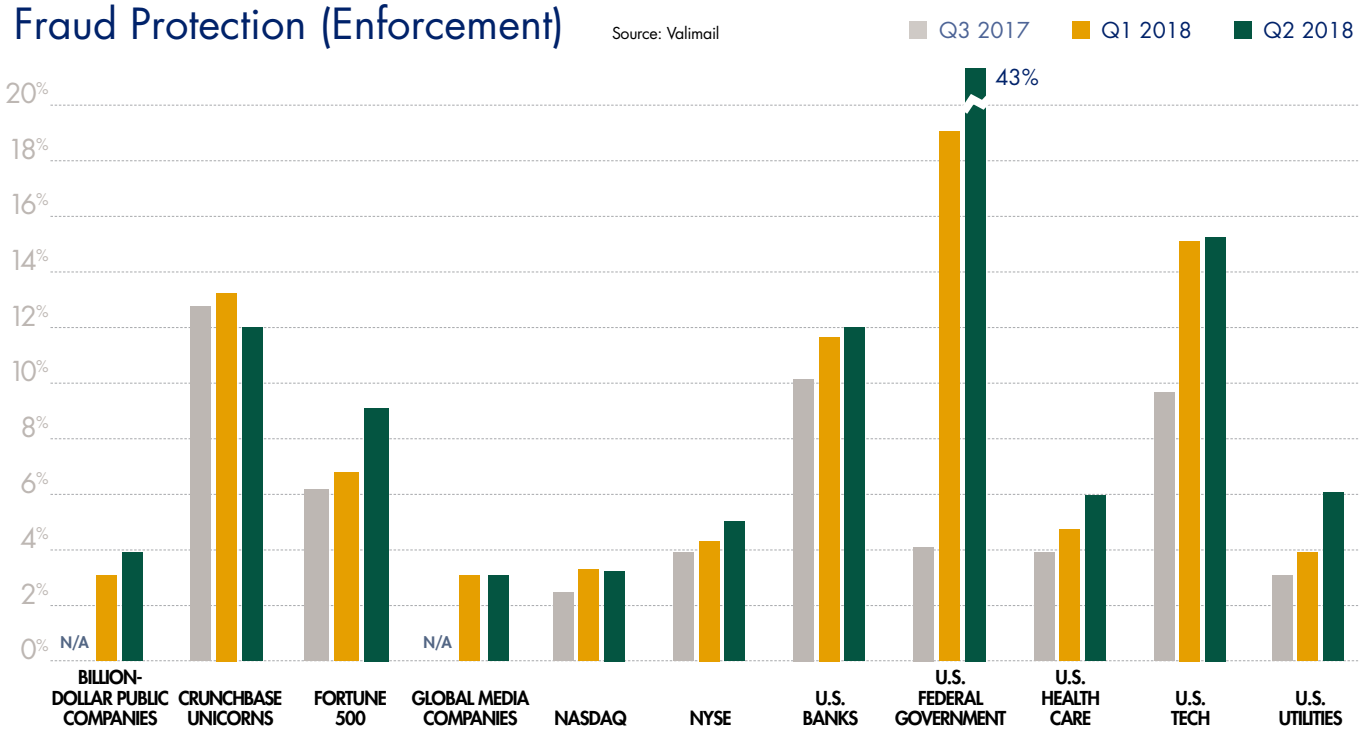


Multiply a category's DMARC usage rate by its enforcement success rate (the inverse of failure rate), and you get its fraud protection rate. In other words, this is the percentage of companies in a given cohort that are protected from fake email by DMARC records that are syntactically and technically valid, and which have been set to an enforcement policy.

Thanks to its high rate of DMARC deployment and high success rate, the U.S. federal government again shows leadership here, with a fraud protection rate of nearly 43 percent. That is a remarkably high figure, and the CIOs and CISOs responsible for this progress deserve congratulations for the progress they have made. There is still a ways to go, of course, as 57 percent of federal domains remain open to impersonation by fake emails.

Other groups showing good numbers are large U.S. tech companies, Crunchbase unicorns, and large U.S. banks: All have greater than 10 percent fraud protection.

There is also cause for optimism among U.S. utilities and U.S. health care companies: Their fraud protection rates have been steadily improving for three quarters.

## SPF Usage

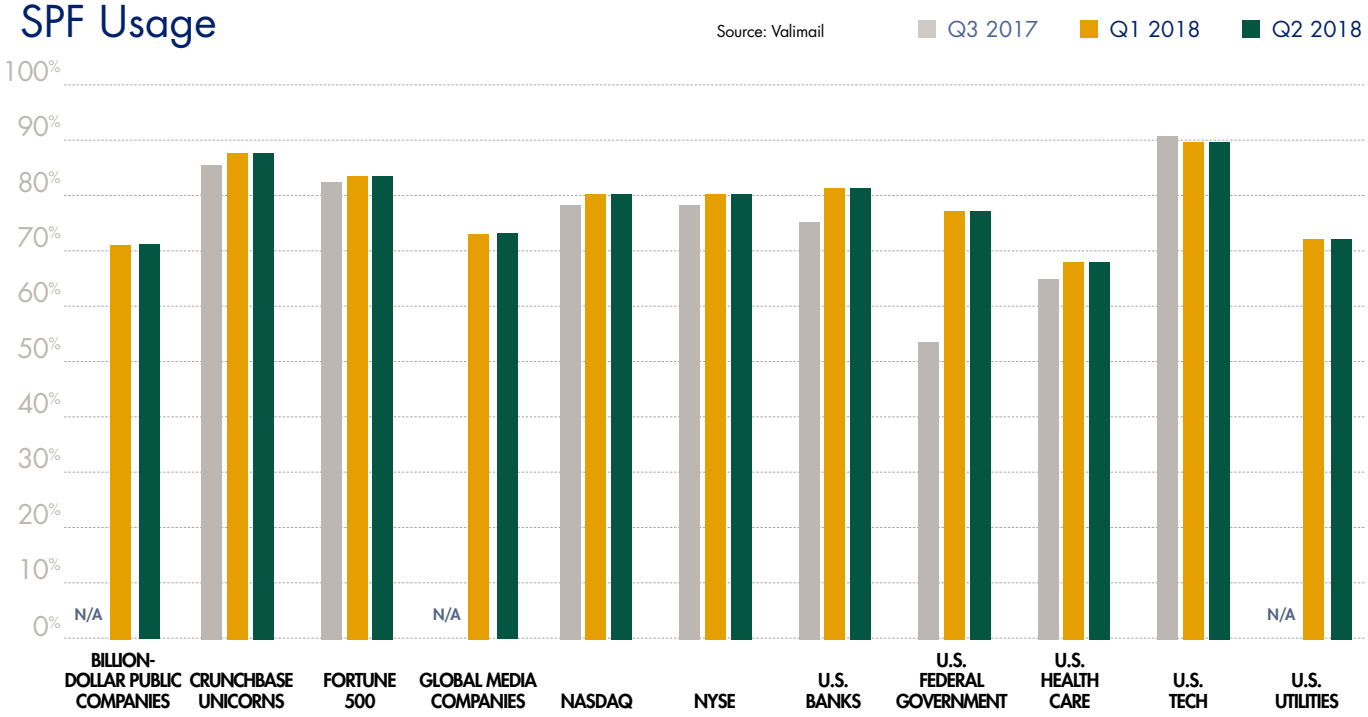


Looking at SPF deployment, it's not surprising to see that the numbers are much higher. SPF has existed as a standard for almost 15 years, and it's widely understood. Because of its value in increasing deliverability for emails, SPF is often recommended as part of email marketing best practices. No industry that we've examined has had an SPF usage rate under 60 percent since Q1 of 2018, and the only category that was lower in Q3 2017 was the federal government.

In all groups, SPF usage continues to grow, albeit slowly. In other words, while this is a mature technology with widespread usage, there is still room for additional adoption — and that's just what we are seeing.

Deploying an SPF record and deploying a correct record are two different things. In nearly every category, we have observed that the number of valid SPF records is 10 to 15 percentage points lower than the overall number of published records. This is due to configuration errors including syntax errors as well as more recalcitrant problems such as the 10-domain lookup limit.[7] Such errors mean that a non-trivial percentage of all SPF records are not going to properly authenticate all the messages that they should authenticate. (And since DMARC is dependent on having correctly configured SPF and/or DKIM, SPF issues often invalidate the DMARC policies that depend on them.)

7 For an explanation of common SPF errors, see "What Is SPF?" on Valimail's website.

## Conclusions

Faced with an avalanche of fake emails, organizations in both the public and private sector are responding by deploying email authentication — with increased rates of DMARC and SPF usage in virtually every category.

While these authentication technologies only stop one type of fake email (the exact-domain spoof, where attackers use a fake From: address), this is a significant step in the right direction as it secures companies (and their customers and partners) from one of the most common — and the hardest to detect — type of fakes.

Driven by the October, 2017 directive from the Department of Homeland Security, the U.S. federal government occupies a substantial leadership position relative to all other groups we studied in its use of anti-impersonation technologies and its success at getting DMARC records to enforcement.

However, there is still clearly a long way to go before the world is protected from fake emails. In future reports we will examine other types of spoofing — and we will continue to report on the world's progress as stopping the crisis of fake emails.

## About the Author

**DYLAN TWENEY | HEAD OF COMMUNICATIONS**
Dylan manages communications for Valimail, including PR, social media, brand awareness , blog content, and research reports.