



MEASURING THE IMPACT OF DMARC'S PART IN PREVENTING BUSINESS EMAIL COMPROMISE

Adam Shostack, Jay Jacobs and Wade Baker

WWW.GLOBALCYBERALLIANCE.ORG

Executive Summary

The deployment of DMARC is an effective way to limit the spoofing of email from protected domains, however, there has been little research on the economic value of deploying DMARC. This paper shows that there is in fact significant value from deployment of DMARC, even looking at only one specific benefit of many (limited to Business Email Compromise, or BEC) and only from deploying DMARC at one of the higher policy levels.

Since June 2016, the Global Cyber Alliance (GCA) has been working to accelerate adoption of DMARC, an email security standard, by providing a set of easy-to-use tools and campaigns to drive deployment. This paper investigates and measures the economic benefit from that work. Having reviewed the available data, we have chosen to focus on Business Email Compromise (BEC) because it is a rapidly growing issue, with high direct losses, and relevant data is available for analysis from multiple sources¹. We derive a conservative minimum bar estimate for the loss avoidance tied to GCA's initiatives and discuss the potential scale of other benefits gained from DMARC.

¹ According to the 2017 FBI/IC3 Internet Crime Report, more losses were attributed to BEC scams than any other type of threat.

In short, for 2018, the estimated value to the 1,046 organizations that have deployed DMARC at a policy level of “reject” or “quarantine,” after using GCA’s tool, is likely \$19M (USD). That assumes:

- **Only 1% of BEC emails that a person receives result in a successful BEC attack.**
- **DMARC stops email impersonation, not other attacks, such as phishing/theft of passwords, identity theft, or other compromises of the computer. Note: these other attacks may be impacted by DMARC, but the pathways are diverse, and measurement is impeded by data quality issues.**
- **There are other benefits from DMARC, including benefits for email deliverability and by deploying DMARC at a monitoring level. Again, there are benefits from both, but these costs are not quantified in this study.**
- **GCA’s public advocacy and training had no other benefits in encouraging the deployment of DMARC.**

This number scales up rapidly as factors are adjusted. For example, if 5% of BEC attacks result in a user taking action, the prevented loss rises to \$66M. These amounts grow over time, so that if:

1. 1,046 organizations maintain DMARC, but GCA's efforts result in no additional deployments (that unlikely assumption makes this a low-water mark); and
2. The economic impact of BEC grows at 5% per year², then the benefit over 5 years would be approximately \$110M. Projecting over a longer timeframe involves increasing uncertainty.

² The FBI released new numbers as this report went to press, showing a 136% increase in BEC from December 2016 to May 2018. Such rates seem unlikely to sustain consistently over a 5 year period.

Moreover, the \$19M estimate is merely from GCA's activity. GCA's other partners, including nonprofit organizations, commercial vendors, governments, and enterprises are also deploying and supporting the deployment of DMARC. For these organizations, the most-affected 1% of them will save \$302,000 per year solely from reducing BEC through DMARC, even assuming only 1% of BEC emails lead to action. If that number grows to 5%, the loss reduction for that 1% is \$1.3M per year. Adding that up across the Internet is a big number, and that is from BEC losses alone. In conclusion, if the people in your company will not be fooled by a targeted email that appears to be from your CEO asking for a transfer of funds to a supplier as an emergency matter, you may not worry about BEC. If, however, there is only a 25% chance your CFO might make a mistake, and you are an average size company, it is more likely than not you will lose somewhere between \$8,800 and \$4,700,000 depending on how lucky you are.



If there is only a **25%** chance your CFO might make a mistake, and you are an average size company, it is more likely than not **you will lose somewhere between \$8,800 and \$4,700,000.**

Contents

Executive Summary	2
1. Introduction	5
1.1 Context and Background	5
1.2 Methodology	6
2. Models.....	8
2.1 Models of Email Attack	8
2.1.1 Business Email Compromise	11
2.2 Data	13
2.2.1 Data for Item 1, Number of BEC Attacks Per Firm Per Year	13
2.2.2 Data for Item 2, Emails Spoofing Domain.....	15
2.2.3 Data for Item 3, Human Action Rate	17
2.3 Other Email Attacks.....	17
2.4 Loss Model.....	17
2.4.1 Loss Per Incident.....	18
2.4.2 Estimating Losses Over Many Organizations	20
3. Benefits From GCA's DMARC Project	22
4. Other Benefits of DMARC.....	24
5. Discussion & Methodology.....	26
5.1 Economics	26
5.2 Phishing.....	29
5.3 Caveats and Limitations	30
5.4 Assumptions.....	30
Conclusion.....	32
DMARC Primer	33
Glossary.....	34
References	35
About the Global Cyber Alliance	38



1. Introduction

1.1 Context and Background

DMARC stands for Domain-based Message Authentication, Reporting and Conformance. It is the simple, trusted, solution that brings together email authentication protocols, and adds reporting and compliance.

There is growing support for DMARC. In June 2016, the U.K. government mandated that all U.K. government departments adopt DMARC, and the EU-CERT has also made a recommendation for the use of DMARC. In October 2017, the U.S. Department of Homeland Security issued Binding Operational Directive 18-01, which requires the adoption of DMARC by federal civilian domains.

GCA discovered, however, that despite the tremendous benefits of DMARC, it was not being widely deployed in the public or private sectors. They wanted to make it easier and affordable for any organization or user to implement DMARC. So GCA created the step-by-step DMARC Setup Guide, currently available in 18 languages, to help organizations of all sizes implement DMARC.

**In August of 2017,
GCA asked for help to:**

- **Measure return on investment from GCA projects, determining the level of risk reduction provided by the resources expended by GCA (and others).**
- **Develop a mechanism for determining systemic cyber risk that can be used for future projects both to measure risk and determine what data needs to be recorded.**

Shostack and Associates and the Cyentia Institute have been helping GCA work towards these goals, with the first project focused on their DMARC work. As we conducted the analysis, it became apparent that the impact of DMARC has both loss-avoidance impacts and brand and deliverability impacts. We have focused this analysis on the BEC subset of loss avoidance, and have developed new mechanisms for understanding that set of losses.

1.2 Methodology

We began with a review of the GCA DMARC Setup Guide tool, and implemented DMARC across several of our domains to gain familiarity with the technology. We reviewed available reports on DMARC drawing heavily but not exclusively on the Cyentia Library, a public and searchable collection which at the time comprised 862 industry research reports curated by the [Cyentia Institute](#). In a first pass, we reviewed 44 matches for the term “BEC”³ and also looked at references to “phish” and “phishing”. In a final review pass, we checked for the same search terms in 300 new documents added to the library to see if additional data had been published. We reached out to GCA partners repeatedly to ensure that we had gathered all the data available. We crafted preliminary models and presented them at the Anti-Phishing Working Group’s E-Crime Conference in San Diego in May, 2018, and then again in a webinar with the goals of getting feedback on the models and data to inform our analysis. Some of this preliminary analysis led us to the understanding of the breadth and diversity of definitions, methods, analyses and production processes associated with published data. We discuss some of these approaches in section 5 to address why we chose not to go down these avenues.

Estimates of economic benefit from DMARC can be grouped into 4 categories:

- **Costs from BEC attacks**
- **Costs from other attacks (phishing, malware and breaches)**
- **Benefits of deliverability**
- **Ecosystem benefits**



We focus on BEC benefits to the organizations deploying DMARC because it represents a major category of threats relevant to DMARC, the data is relatively closely grouped, with much of it coming from the FBI⁴. The FBI remains a credible source of crime data. This methodology provides a floor, rather than a complete estimate. The other types of benefit, with their higher variance, may be addressed in future work.

We consider these benefits to accrue to those who have defined a DMARC policy of ‘quarantine’ or ‘reject.’ (Organizations can also deploy DMARC with a policy of ‘none,’ see section 4 for a discussion of the benefits from the none setting, and section 7 for more on the effects of these settings.)

³ More precisely, filter(token0 == “bec” | token0 == “business” & token == “email” & (token2 %in% c(“compromise”, “spoof”, “spoofing”)))

⁴ As this report went to press, the FBI’s Internet Crime Complaint Center released new data, which we have not incorporated. (FBI18) They note that ‘between December 2016 and May 2018, there was a 136% increase in identified global exposed losses,’ which means that our estimates are again, a low-water mark. (EDIT-SA-19)

Reviewers pointed out that the per-organization loss avoidance numbers “don’t seem like they’re enough to drive activity.” We are reminded of a story about Miss Manners, who gets into a taxi at JFK and asks to be taken to Newark. “Newark?!?” the driver exclaims, “But that’s so far!” “Very well,” Miss Manners responds, “where would you like to go?” We have gone where the data has taken us, but would add that the cost of deploying DMARC does not have to be particularly high. One author of this report did it over the course of a few hours. Large organizations may take months to track down all of the various departments, contractors, and service bureaus sending email on their behalf around the globe. That time may be because finding email senders can be tricky, and because some emails may be sent monthly or even quarterly by different processes. The work to get all that “under control” is also a hard-to-measure benefit.

We also chose not to address estimates of criminal gain from these activities, as criminals are often untrustworthy. We chose to not address societal losses from these activities.

The rest of this paper is organized as follows. Section 2 presents the models of email attacks and BEC that we used, the data we considered, and the loss model on which we settled. Section 3 presents the results of a set of simulations based on the data and loss model. Section 4 presents an overview of the other benefits of DMARC. This is followed by a discussion of the methodology (section 5), some conclusions (section 6), a primer on DMARC (section 7), a glossary (section 8), references (section 9) and finally, information about the authors.

Recommendations

- **Use the Global Cyber Alliance DMARC Setup Guide at <https://www.globalcyberalliance.org/dmarc/> to check the state of your organization’s DMARC records today.**
- **If you do not publish DMARC records, set them up and start gathering data. DMARC allows you to do this without impact to your email infrastructure or deliverability by setting a “none” policy.**
- **Ensure that your inbound email is configured to respect other organizations’ DMARC policies.**



2. Models

We use a variety of models, including models of email attacks and models of losses, to help us understand the impact of DMARC on inbound BEC. We present models and then the data which informs their use. We look at the email models first, then the economic models.

2.1 Models of Email Attack

To assess the impact of DMARC, we consider the likelihood that DMARC prevents an inbound email from reaching an inbox. DMARC actually makes spoofing email from an organization harder, and thus applies to all mail, and we are looking at this subset for methodological reasons. We present a modified attack tree⁵ for understanding the flow of inbound email, which is important in understanding how DMARC helps. We then compare this attack tree to Agari's email threats taxonomy.



Dmarc actually makes spoofing email from an organization harder.

Our phishing attack tree is scoped to email attacks on a target that the attacker has not yet compromised. It considers a message progressing through time and systems on the vertical axis. A message starts when an attacker sends it, goes to the target's mail server where it is subject to DMARC testing, then the target's inbox where it can be acted upon.

⁵ A threat tree is a structured technique to enumerate all the paths by which a system may be attacked. Our tree differs from a standard threat tree in that it shows time and systems, and not all nodes are threats. It is also not a complete enumeration of all threats. See (Amaroso) for a detailed description of threat trees.




Figure 1

An attack tree model of phishing

Safer states are in yellow, human actions in orange, and incidents in red. The “passes by DMARC” state is a more dangerous state. We refer to “passes by DMARC” to mean that the message is not blocked when the sender has published a DMARC policy of “quarantine” or “reject”. We treat human action as the final step; that is, no controls intercede after a person has acted. This is a simplification that makes our modeling simpler. The exaggeratory effect can be tempered, and the other controls simulated, by wrapping them into the rate of human action on messages, which is a parameter in our models.

There are responses to a BEC email other than replying, but usually there’s a cycle of reply before a bank transfer is initiated.

If the victim is at bank.com, then the attacks would be:

Attack	displayed	Really
Spoof	Alice@bank.com	Alice@attacker.com
Look-alike	Alice@b4nk.com	Alice@b4nk.com
Display Name	Vice President Alice	Vice President Alice <alice@attacker.com>

INBOX IS A CRUCIAL STAGE

We focus on email delivered to the inbox. We focus on this key stage because an email which is delivered is potentially harmful, while one whose delivery is prevented is a denial of service on email infrastructure and has only a capacity-planning impact.

We choose to not distinguish between inbox and “spam box” because implementations vary. Most people have effectively been trained that they need to look in their spam boxes for missing messages. That training persists even as some email providers have become very accurate. The training persists because people unlearn behaviors slowly, many people use many mail systems with inconsistent filter quality, and they are trained to expect that emails will be filtered.

This training comes via messages like the following (which were recently received by the authors):

- **“Please add us to your address book or safe list.”**
- **“To ensure delivery of future emails, please add...”**
- **“Please add [redacted] to your address book.”**
- **“To make sure you continue receiving my newsletters, please add my email to...”**

The technically inclined may note that these taglines decorate emails from bulk senders, a distinction most people will not bother to consider or reach. This failure to distinguish could lead us to only look at DMARC “reject” policies, but in practice, DMARC “quarantine” is used as one of many input signals to spam detection, and the worst spam doesn’t even make it to a spam box.

There are controls in place that work after an email reaches the inbox. These include URL rewriters, anti-malware, and business/policy controls on funds transfers. These controls are diverse and complex, and appropriate use of DMARC protects against failures in these controls.

The controls are diverse and complex, and appropriate use of dmarc protects against failures in these controls.

2.1.1 Business Email Compromise

Our model of DMARC impact is focused on Business Email Compromise (BEC). This focus is driven by both the reliability of the statistics from the FBI/IC3 and the fact that reported losses from BEC are greater than the next 6 types, combined.⁶ (FBI17)⁷ Building a model focused on BEC will enable the use of specific data sources and capture a major component of losses caused by email-based attacks. Once we have a base model based on BEC, we can extend estimates into other types of attacks.

Building a model of BEC can be done with the following steps:

1. **Build a distribution for the number of BEC attempts received by organizations annually.**
2. **Establish the proportion of BEC attempts that use domain spoofing as a tactic.⁸**
3. **Establish the proportion of attempts prevented/blocked by the implementation of DMARC in a blocking mode (“quarantine” or “reject”) to start with estimations of BEC attempts per year.**
4. **Establish the probability of human action resulting in a funds transfer. Human action refers to all the steps from receiving an email through approving funds transfer. Thus human action here is distinct from the click rate on a phishing link or malicious attachment.**
5. **Establish a distribution of losses from BEC attacks. As previously mentioned, losses from BEC are dramatically higher than other types of cyber incidents.**
6. **Determine the loss avoidance, based on steps 2 and 4, and by comparing those with DMARC to those without.**


BEC scams have accounted for \$12.5 billion dollars in losses around the world the last five years.

⁶ Losses from BEC are 3.2x “Confidence Fraud/Romance, 4.7x non-payment, 7x investment, 8.8x personal data breach, 10x identity theft, and 11x corporate data breaches (reported on by the FBI). Together, these 6 amount to \$654,230,043. (EDIT-VM15)


⁷ The wide variety of “headline numbers” around intrusions caused disagreement, even within the team, on relative magnitudes. There are reasons to think the FBI gets better numbers around BEC than around other computer crimes, including the use of police reports for insurance, the ease of measuring BEC costs compared to costs of other types of computer crime, and the desire of many organizations to sweep data breaches under the rug.

⁸ Domain spoofing means that someone sends email claiming it’s from alice@bank.com, when it’s really from alice@...

We can tie this list of steps to the attack tree as follows:



Another way to represent this funnel is to consider what we can see and what we cannot as the attack flows through from attack to financial loss. In this figure, based on (Romanosky), time flows left to right. Observed data is in filled circles; data that is not observed/used in this analysis is represented by empty squares.



2.2 Data


2.2.1 DATA FOR ITEM 1, NUMBER OF BEC ATTACKS PER FIRM PER YEAR

We found several sources discussing this statistic, and it provides a fairly good launch point for the model.


- **"On average a targeted organization has 5.2 BEC emails sent to them each month."** (SYMC)
(Note: At 5.2 per month, the annual average is about 62 BEC emails.)
- **"On average, companies were targeted by 18.5 fraudulent emails per quarter."** (PROOF17)
(Note: at 18.5 per quarter, the annual average is about 74.)
- **"In the second half of 2017, BEC attacks continued to accelerate with 96% of organizations analyzed by Agari being attacked at least one time, and with the average business experiencing 45 BEC attacks from June through December 2017."** (AGARI01)
- **"In the first three months of this year (2017), nearly 85% of organizations were targeted by at least one BEC message."** (PROOF01)

While it may seem that the annual average fluctuates quite a bit here (62, 74 and 45 being cited) this does help us narrow down our distribution. Another interesting aspect is the wording chosen by two different sources about the proportion of companies receiving "at least one" BEC attempt. This leads to an assumption that the majority of organizations will see very few attempts annually, but a minority may receive many attempts in a given year. This is an assumption in how we created our distribution, but it is in line with other count/event data (such as the number of deaths from [horse kicks in the Prussian Army](#)).

Using the above as input, we created the following distribution for the number of BEC attempts annually. Note that the next graph "BEC Emails Per Org" starts at 1, not zero.



(Same Data, just a different view, bars are binned by 10 emails)



This distribution (using 100,000 simulated organizations) has the following statistical properties:

- **95% of orgs saw at least 1 annually**
- **52% saw 10 or fewer BEC attempts annually (including 5.2% who saw 0)**
- **Average⁹ across all orgs is 60**
- **17.4% of orgs saw more than 50 BEC attempts annually (conversely, 82.6% saw 50 or less)**

Note that these properties are very consistent with the data sources cited above.

⁹ Losses from BEC are 3.2x "Confidence Fraud/Romance, 4.7x non-payment, 7x investment, 8.8x personal data breach, 10x identity theft, and 11x corporate data breaches (reported on by the FBI). Together, these 6 amount to \$654,230,043. (EDIT-VM15)


2.2.2 DATA FOR ITEM 2, EMAILS SPOOFING DOMAIN

Since DMARC (in enforcement mode) will prevent the delivery of emails that attempt to spoof the target domain, if we estimate the proportion of BEC attempts that use that technique, we can effectively identify the proportion of attempts that will be blocked by DMARC. The published research represents distributions with a single representative number, and those are below with our emphasis added:

- **GreatHorn reports out of 537,617 spear-phishing threats, 44,726 were “direct spoofs”, which is about 8% of the fraudulent emails they studied. (GREATHORN)**
- **Agari covers this statistic quite thoroughly (AGARI01):**
 - “...Agari’s research shows that 12% of BEC attacks use spoofing; 7% a combination of look alike domains and display name deception; and 81% pure display name deception.”
 - “For organizations that use Proofpoint, 95% of attacks that went undetected used display name deception and 5% domain spoofing.”
 - “For organizations that used Microsoft EOP with no third party SEG, 90% of attacks were display name deception, 7% domain spoofs and 3% look alike domain based BEC”
 - “For organizations that use Google G Suite with no third party SEG, 93% of the attacks that were not blocked used display name deception and 7% domain spoofing.”
 - “For small businesses, 90% of BEC attacks observed used display name deception, 6% used domain spoofing and 4% used look alike domains.
 - “For medium businesses, 95% of BEC attacks used display name deception, 3% used domain spoofing and 2% used look alike Domains.
 - “For large businesses, 75% of BEC attacks used display name deception, 16% domain spoofing and 9% look alike domains.”

While the average percent of inbound emails spoofing the recipient domain is a single number, we don’t know what it is. However, using the published measurements (listed above), we can capture our uncertainty around that value and create a distribution of probable values. The result (visualized below) represents our understanding and uncertainty around the percent of inbound emails spoofing the recipient domain. This is also an estimate for the proportion of BEC attempts blocked by DMARC, if configured to do so.


■ Probability of Inbound Email Spoofing Recipient Domain



This distribution has the following properties:

- ➔ Average is 9.7% (arithmetic mean), minimum of 3%, maximum of 19.8%
- ➔ 50% of the distribution is between 7% and 11.1% (highest density interval, 50% mass)
- ➔ 90% of the distribution is between 4.9% and 14.4% (highest density interval, 90% mass)

■ Derived Metric: BEC Emails Blocked by DMARC



By simulating the number of BEC emails received by 100,000 organizations and the probability of BEC emails spoofing the target domain, we can estimate the expected number of blocked BEC attempts annually. Combining this metric with the estimated losses from successful BEC attacks, we can calculate the potential savings by implementing DMARC for an organization.

2.2.3 DATA FOR ITEM 3, HUMAN ACTION RATE

The human action rate is a measurement of the funds transfer actions resulting from a BEC attempt that lands in an inbox. We had a challenge in finding any reliable data around the success rate of BEC attempts. So we referenced the success of phishing as a possible proxy measurement and set up “reference points” for the rate of human action. We will simulate results if the rate of human action is 25%, 10%, 5% and 1%. There are reports of phishing click rates which are higher (Wombat/Proofpoint 2018), but we do not believe those rates translate into BEC successes.

2.3 Other Email Attacks

We believe that the data available to us on BEC attacks is higher quality (people are less likely to exaggerate on a police report). There are costs to organizations from cleaning up intrusions, password resets, or from data breaches, losses from phishing to both banks (who may bear the cash costs) and consumers (who must spend time and may fear an accusation of fraud), and to everyone from the need to invest in security measures. Good overviews of these are in (Anderson, 2013) and (Riek, 2016).

However, that the data is higher quality should not be taken for a claim that the data is complete. Victims of BEC may not report it because the loss is small, they lack resources to support a criminal investigation, or they are concerned about involving law enforcement because of immigration status or other reasons. Victims of BEC outside the U.S., or whose bank is outside the U.S. might not report it to the FBI, yet they do in surprising numbers. The FBI report has an infographic of the top 20 foreign countries by victim, and each of Canada, India, the U.K., Australia, Mexico, the Russian Federation and Brazil have reported over 500 crimes to the FBI/IC3, and between them they have 10,139 reports. (FBI17, page 18)

2.4 Loss Model

For costs, we start from a standard model of adverse events, that they have a frequency and an impact. Losses are dependent on the frequency of attacks and the frequency of success.

“If, however, there is only 25% chance your CFO might make a mistake, and you are an average size company, it is more likely than not you will lose somewhere between \$8,800 and \$4,700,000, depending on how lucky you are.”

2.4.1 LOSS PER INCIDENT

When formulating the loss distribution, we focus on BEC since the losses recorded from BEC completely dwarf other types of losses (FBI17). In modeling BEC losses, we have a few data points to go from. First is the statistics from the FBI/IC3 which stated, “In 2017, the IC3 received 15,690 BEC/EAC complaints with adjusted losses of over \$675 million.” (FBI17) From this we can calculate the arithmetic mean at just above \$43,000 per loss. Other FBI statistics (IC3 BEC) show \$131K in average loss ¹⁰ for U.S. victims, and a \$71K worldwide average loss. And recently updated statistics from the FBI (FBI18) shows over \$12 billion in losses over 78,617 domestic and international incidents, which averages out to about \$160,000 per organization. With these measurements, differing by geographics and timespan, we know the estimates of the arithmetic mean of losses from BEC to be around \$43K, \$71K, \$131K and \$160K. While the range and variance in those estimates may seem to vary wildly, keep in mind they are the arithmetic mean which is highly susceptible to extreme values. And looking across the headlines, we find several examples of very extreme losses:

- **“Leoni AG, the fourth largest wire and cable manufacturer in the world, became a victim of a BEC attack when its Chief Financial Officer (CFO) was tricked into transferring about US \$44.6 million to a foreign account.” (TrendMicro)**
- **“Online wire transfer provider Xoom was probably the year’s first victim of a Business Email Compromise (BEC) to the tune of \$31 million.” (Verizon16)**
- **“...Belgian bank Crelan announced they were the victim of a €78 million BEC fraud.” (about \$91 million USD) (Verizon17)**

While there are several more examples like these, the “headline” losses are consistently in the tens of millions. How do we reconcile the millions in losses in the headlines with the data from the FBI/IC3 statistics (\$43K - \$160K)? The answer lies in the shape and distribution of losses, and by understanding that, we can arrive at a descriptive distribution of losses for BEC that goes beyond any single estimation of an average. The headline cases are headlines for a very specific reason: they are rare and they are extreme. So we know two facts about the distribution: 1) the mean is somewhere between \$43K and \$160K, and 2) the extreme but rare cases go into the tens of millions. Losses from BEC are clearly a skewed and long-tailed distribution.


As a proxy to estimating the distribution, we can turn to other loss distributions, and since some of the authors have experience with loss data (Verizon15), we know the distribution of losses is roughly lognormal in nature. This means if we plot the losses on a log scale, it will appear like a normal (bell-shaped) distribution, with the peak around the median. Calculating the arithmetic mean (average) on a distribution like this will arrive at a value that appears off-center when visualized. This happens because lognormal distributions have an extremely long tail and the events on the tail (e.g. the “mega-breaches” that make the headlines) have a large influence on the mean, making the mean much higher than the median.

¹⁰ The IC3 refers to “exposed dollar loss,” which “includes actual and attempted loss.” (IC3 BEC, their footnote 3). This definition may appear on first blush to be overly broad, but it actually aligns well with our measure, which is focused on email delivered to the inbox.

As an example of that effect, the NetDiligence Cyber Claims study looks at insurance claims and reported, “The average (breach) cost for the period 2014–2017 was \$394K; the median cost was \$56K.” Additionally the research stated, “The smallest breach cost reported was \$110 while the largest was \$16.8M.” That large separation between the median (\$56K), the mean/average (\$394K) and the maximum (\$16.8M) establishes that losses are long-tailed and cannot be modeled with a simple (linear) distribution.

We have created a distribution to represent our estimates of U.S. losses from BEC/EAC events, which is shown in the Figure of that label. Our estimated distribution lines up with the observed data points around the reported averages and encompasses the “headline” events, and it is in line with previously studied loss distributions (Verizon15). To generate this distribution, we leveraged the betaPERT distribution and forced a heavy-tailed distribution by using the natural log of the parameters and scaling the distribution back up by applying the exponential function. The results are shown in Figure “Estimated Distribution of U.S. Losses from BEC/EAC Events”. The solid line represents the arithmetic mean (\$108,000). It’s interesting to note that the majority of BEC losses would have to be small for this distribution to be accurate.


■ Estimated Distribution of Losses From BEC Attacks



Estimated Distribution of U.S. Losses from BEC/EAC Events

2.4.2 ESTIMATING LOSSES OVER MANY ORGANIZATIONS

We can combine all of the estimations we've created so far to get an overall estimation for the financial effect of DMARC on a single organization and then extrapolate back out to the financial effect for the collection of organizations assisted by GCA. We represent financial effect as "loss avoidance" by simulating organizations with and without DMARC as a security control. The difference between the losses with and without DMARC represent how much loss would be avoided with the implementation of DMARC.



The way to read the above plot is to start with a rate of human action, if we assume that about 5% of BEC attempts result in human action (the orange line marked "5% act"¹¹). All the way to the left, we see about 20% of organizations will not avoid any losses (either because they are not attacked or because the attack uses a vector that DMARC cannot stop), 15% of organizations will avoid at least around \$7,000, 10% of organizations will avoid about \$30,000, and so on. For convenience, we selected a few points along the curves for the table below.

¹¹ We avoid the acronym HA for human action to avoid implying any contempt for people who handle lots of email daily and sometimes make mistakes in doing so.

■ Table of Estimated Losses Avoided (Per Org.) from DMARC Stopping BEC Attacks

Percentage of Organizations	Minimum Loss (USD) Avoidance from DMARC implementation			
	25% ACT	10% ACT	5% ACT	1% ACT
1%	4.7M	2.2M	1.3M	302K
2%	2.5M	1.1M	583M	118K
5%	854M	333M	145M	20K
10%	277M	88M	29M	1.2K
20%	49M	10M	0	0
30%	8.8M	0	0	0

With estimates about the loss avoidance for a single organization, we will extrapolate across the organizations assisted by GCA in section 4.




3. Benefits From GCA's DMARC Project

If we extend the per organization result to the 1,046 domains that have leveraged GCA to implement DMARC as of June 2018, we can get an estimate of the direct loss avoidance GCA will help companies realize over the next 12 months. From the previous outcome, we can simulate the combined effect on 1,046 organizations for a variety of human action rates. The subsequent charts show our best effort estimation as a point with the margin of error represented by the line behind the points.¹² For example, if we assume the action rate is 1% (again, meaning each attempt has about a 1% chance of resulting in a completed action by a person in the target organization), the average sum of BEC losses avoided per year for 1,046 organizations assisted by GCA is about \$19 million, with a 95% credibility interval from \$7.3 million to \$55 million.

We consider avoided losses over 12 months starting in June 2018 because domains have turned on protection at various times over the preceding years and this is a simpler model. We do not expect the costs of an individual BEC event to decrease substantially over the next 12 months.

As a reminder, we simulate 100,000 organizations per year, and we simulate the frequency that their employees click leading to BEC losses. We then plot that data for each click rate into Figure "GCA-Enabled Loss Avoidance per Year (USD)." The center line represents the most likely outcome, and the shaded region expresses the 95% credible interval around the most likely outcome. In reading this, first pick out an estimated Human Action Rate and look across the credible interval (shaded) and the most likely outcome (solid line). For example, at 5% action rate, the most likely outcome is about \$66M, but given the variability in our data, the loss avoidance could be between \$36M to \$118M (the shaded region).

¹² Check out that sleight of hand! We go from "domains" to "organizations" as if they're the same thing. Also, pay no attention to the organization behind the domain curtain, but do pay attention to assumption #7.




ACTION RATE	LOWER CREDIBLE INTERVAL (95%)	MOST LIKELY VALUE	UPPER CREDIBLE INTERVAL (95%)
25% ACT	25% ACT	10% ACT	5% ACT
1%	4.7M	2.2M	1.3M
2%	2.5M	1.1M	583M
5%	854M	333M	145M
10%	277M	88M	29M
20%	49M	10M	0

Assuming that BEC remains a relatively easy crime to engage in, with low barriers to entry and high payoffs, we can expect that next year, the same benefit will accrue to those who have turned on DMARC, and additional benefit will accrue to those who turn it on, and so benefits next year should be greater.

4. Other Benefits of DMARC

There are benefits from DMARC even at a policy of “none”. These benefits are visibility into email sending infrastructure and domain spoofing which may be taking place. We do not attempt to quantify these. The authors were surprised to see the spoofs on one of our little domains when we first turned on DMARC. False email volume, shown in red, dwarfed real email, shown as the little tiny green blotches at the X axis. (Visualization of shostack.org email, courtesy of dmarcian).



There are security benefits to others in turning on DMARC. The drop off in February was probably when we moved to DMARC “quarantine”. Our customers, partners and friends became better protected against email claiming to be from us. Security awareness courses often focus on “people you don’t know” and spoofing of emails from lawyers, accountants and other service providers has been a fruitful avenue for attackers.

DMARC has non-security benefits, such as increased email deliverability and brand protection. These may be quite large. Email marketing may be a multi-billion dollar business. (Statista) claims that it will surpass 3 billion USD in 2019, and (Illumenmedia) claims a return of 44x on email marketing spending. If we accept that number and assume that DMARC increases deliverability by a few percentage points, then the return might go from 44x to 46x, which would globally represent a \$6 billion gain. Such measurements are difficult because we lack data on deliverability, expertise such as the ROI deltas, and their distribution. The problem is also complex because we would expect that businesses that stand to gain the most will have already invested in DMARC and other deliverability tools. It turns out that this is not the case. For example, Valimail reports that as of Q1, 2018, less than 40% of U.S. banks are using DMARC. We expected that number to be substantially higher. (Valimail).

As noted in the introduction, there is also organizational transparency and visibility to the deployment of DMARC. Knowing who sends email to the customers of a large organization is, as it turns out, a non-trivial task. When email was cheap, barriers to sending emails low, and GDPR didn't exist, it was easy to set up email to stay in touch in a variety of ways. A DMARC policy of reporting without enforcement can illuminate the many organizational centers who have "always done things that way."

Also related to the measurement of deliverability benefit is the challenge of going from a list of domains to an understanding of the owners of those domains, and if there are businesses behind them. For example, a domain may be parked; it may be for use by a non-commercial entity. It is also challenging to go from the subset of businesses to an understanding of if or how each might benefit from email for sales, marketing or other purposes. Instead, we assume that the distribution is roughly logarithmic with a median of X dollars. It seems reasonable to expect that larger businesses that are highly dependent on email will have already spent time on deliverability.

Some portion of DMARC benefits are likely gated on ecosystem improvements. When everyone turns on DMARC, it is easier to block non-DMARC email. These "cooperation benefits" are not quite identical to a Nash Equilibria, because each party individually has motivation to improve, and there are benefits to others moving.

Case Study - Aetna

Each year, DMARC prevents approximately 60 million fraudulent email messages from being delivered.

The result is lower risk and better engagement: click-through rates on Aetna emails improve by 10 percent every year. DMARC adds a trust component to emails and is a core component to Aetna's trusted email program.


5. Discussion & Methodology

5.1 Economics

We started by reviewing the literature and studying existing models of phishing. The Agari taxonomy came close to our needs but does not explicitly tie actions to losses.

We reviewed a set of economic models of cybercrime, including the U.K. government's (HomeOffice) and (Anderson12). We note that our numbers are what they categorize as direct losses (in a spectrum of direct losses, indirect losses and protection costs). The GCA organizations are not categorized geographically, while the Home Office numbers and part of Anderson are, which makes direct comparison difficult. The same geographic issue makes a direct comparison with (Riek) difficult. Anderson and colleagues do offer global estimates, but their work was done in 2012 before the rise of BEC.

Our past experience includes use of rich datasets, particularly the Advisen breach data gathered from court filings, and authors Baker and Jacobs were principals behind the Verizon Data Breach Report.



Source: [Phishlabs](#)

The loss models we considered were:

Method 1

“Cost per record” is a model of breach losses based on cost-per-record, as there are well-known numbers associated with that metric. Unfortunately, they are not fit for our purposes for many reasons, including:

- **Assumption that costs increase linearly with records lost. This is not the case. Small breaches spread large fixed costs over fewer records (very high per record cost), while large ones gain economies of scale for response costs (very low per record cost).**
- **The data is heteroskedastic, meaning the variation is not constant. There is more variation in large breaches than small ones.**
- **Not all incidents are breaches.**
- **Phishing/BEC losses may have outcomes that are not “record-centric breaches”. (How many records are compromised in a typical BEC?)**
- **Phishing/BEC losses may have outcomes that are not “record-centric breaches”. (How many records are compromised in a typical BEC?)**

Method 2

“Cost Relates to Organization” is based on the expectation that organizational factors such as size, revenue, number of employees or sector are controlling factors in the cost of an incident. Prior work by Jacobs (Verizon15) stated, “Breaches with equivalent record loss had similar total costs, independent of organizational size.” Which indicates that these organizational factors contribute weakly to incident costs.

Method 3

“Cost Relates to Data” is based on the expectation that the cost of an incident relates to the type of incident, and so, for example, a phishing incident costs more than a lost backup tape. Unfortunately, we are not aware of a rich dataset that includes root causes across a majority of its data points, and so using “cost relates to data” would involve throwing away a great deal of information.

Method 4

“Nuanced Cost Distributions” is based on the frequently expressed hope that each incident is different, and that if we captured enough data points on past incident we could be more predictive about costs. Unfortunately, such a dataset is expensive to construct and using it predictively would require capturing the same number of data points per possible victim. (If we could capture less, then implicitly, the other data points aren’t important.) That capture would be expensive. Imagine for a moment that we want to say that organizations that rely on email for sales get more value from DMARC than those who use a website for sales. That would require discerning the primary sales model for tens of thousands of organizations. Even if we could reliably do that in 5 minutes per domain, for 10,000 domains it would be 833 hours to capture that single data point.

Method 5

“Simple Distribution” Given the issues enumerated above, we plan to use a simple distribution derived from available data points with references to other work studying observed distributions of losses from cyber events.

5.2 Phishing

We consider this likelihood in a static model, where attackers do not alter their behavior in response to defender behavior. This is reasonable because attackers, working in bulk, alter their behavior relatively slowly, as shown by inbound emails that are not authenticated via DMARC. A more dynamic model may be desirable, but data to support it is not available to us. In particular, a more dynamic model requires that we understand either the behavior of a single spammer relative to DMARC being turned on at a domain, or the emails received by a representative set of domains as they activate DMARC. We would want to understand if the domain spoofed emails simply disappear, or if they were replaced with emails with visually similar domains, friendly-name spoofing.

It would be great to know what fraction of phishing attacks (or successful phishing attacks) lead to major incidents, but that number is difficult to pin down. The 2017 Data Breach Investigations Report (DBIR) from Verizon (Verizon17) estimates approximately 10% of major incidents reported came from phishing. The 2018 DBIR estimates 1,192 out of 30,362 (4%) of incidents and 236 out of 1,799 (13%) of data breaches involved phishing techniques (Verizon18). This is a funnel of “phishing attacks” of which some fraction are “successful attacks”. Of successful attacks, some fraction become “incidents” (others are not incidents because an employee’s bank account is compromised via their browser at work, but that’s an incident for the bank, not the employer). Some fraction of incidents become major incidents with costs that are aligned with the Advisen data.

Our attack tree relates closely to the sender portion of Agari’s Email Threats Taxonomy (Agari) but is more focused on the problem DMARC solves. Agari’s is more comprehensive, contains details not relevant to this analysis, and does not show time or action. BEC is a subset of cons, which also include “lost passport” and “Nigerian Prince” scams.

Agari’s taxonomy makes clear the difference between various forms of imposter activity: spoofing, look-alike domains, and display name deception. DMARC only has an effect on spoofing (of a domain), and so that is the focus of this work.

5.3 Caveats and Limitations

Because moving to DMARC “quarantine” or “reject” may seem like a “scary” change, many of the domains which were entered into the tool have not made a change to DMARC “quarantine” or “reject”. The data is sensitive to the results of the monte carlo simulations; a long tail event can change the data.

Of the remaining N domains, it may be that those domains are considering making their DMARC policies more strict.

5.4 Assumptions

To perform any analysis, we must make assumptions. Some of the important ones that we have made are assembled here. We believe these are defensible.

1. **We use a lower human action number for BEC email to transfer than is commonly reported for clicks in phishing/malicious attachment email. We think that this is reasonable because there are several steps between initial contact and funds transfer.**
2. **The FBI reports of BEC are an undercount because some fraction of victims will say “but what will the police be able to do,” and some smaller number may be insurance fraud, netting out to the FBI reporting a smaller number of events than really happened. We debated the correction factor for this. We think it’s reasonable to think that most organizations report to the police because a police report is a step to getting an insurance payout. We have also personally experienced calling the FBI because of a fraud incident, and the FBI didn’t want to take a report of an incident under \$100,000. That matches our prior expectations. But the IC3 doesn’t discriminate in the reports it accepts. Also, IC3 is not as well known as the FBI as a whole. FBI data is largely US-centric. (It may not be US-specific because multinational corporations may report to the FBI for various reasons.) This number is highly impactful on the model.**
3. **The frequency of BEC being sent or acted on has not changed since the FBI issued the two 2017 reports on which we rely. There is likely more BEC email because of the ease of engaging in this crime. It may be acted on at a lower rate because of publicity, or at a higher rate as attacker tactics evolve.**
4. **The cost of a BEC incident has not changed since the FBI issued the two 2017 reports on which we rely.**
5. **People notice email frauds (friendly name, visually-similar domains, spoof domains) at similar rates, and choose to discard the email once they have noticed anything questionable. This discounts the effectiveness of advice like “carefully check the email address”.**

6. The probability of BEC compromises are independent; that is, after a first compromise a second one is equally likely. We think it's likely that the odds of being successfully compromised by BEC decline after a first success, but do not model that because the rates of compromise are low, and it reduces the parallelism of our models and thus overall comprehensibility.
7. Each organization uses GCA to assess a single domain, and thus "domains", which are what GCA tracks, are equivalent in number to organizations. This simplifying assumption is likely an overcount of organizations, as many organizations have many domains, and likely an undercount because many of those domains are not used to send email. We assume those two biases balance out, although a reviewer raised an interesting question as to whether or not it matters that the domains are not usually used to send email. Perhaps the spoof would work regardless. We have not explored this fully because we lack data.
8. The mail servers for organizations implementing DMARC to protect their outbound email are configured to respect DMARC for inbound email.
9. We ignore a set of scenarios where an organization is already compromised in some way, including that a server or account/credentials are already compromised.
10. BEC attacks conform to a Poisson distribution.
11. The FBI categorization of BEC is broader than funds transfer.¹³ Because the FBI does not break out BEC subtypes, we treat all FBI-reported BEC as funds transfer.

¹³ (FBI17), page 12: "BECs may not always be associated with a request for transfer of funds. In 2016, the scam evolved to include the compromise of legitimate business email accounts and fraudulent requests for Personally Identifiable Information or Wage and Tax Statements commonly known as W-2 forms for employees. In 2017, the real estate sector was heavily targeted with many victims reporting losses during real estate transactions."

Conclusion

BEC is an expensive threat, and the work that GCA has done to date has likely led to annual savings of \$19M. As more organizations take advantage of their work, and as time passes, those avoided losses will accumulate to the good of all organizations that rely on DMARC.

In addition to the growth in impacted organizations and the accumulation of value over time, BEC is also a rapidly growing threat, and it bypasses controls which filter on URLs or attachments. Deployment of DMARC can help reduce the odds of costly losses. The GCA work helps reduce overall losses from BEC, and also helps executives understand how their domains are being abused without a lot of complexity. While there may be complexities, an easy first step is to see if your organization gets 3 green checkmarks at <https://dmarc.globalcyberalliance.org/>.

DMARC Primer

DMARC is a security standard which allows a domain to declare a policy of how it wants email processors to examine and treat emails claiming to be from that declaring domain. The core policies are “none” (do nothing special, and optionally report), or “quarantine” or “reject” messages which fail various other security checks.

dmarc.org

<https://dmarc.globalcyberalliance.org/about-dmarc/>

https://www.valimail.com/resources/white_paper/what-is-dmarc/ https://www.valimail.com/resources/white_paper/operationalizing-email-authentication-a-systematic-approach-to-email-authentication/

<https://cyber.dhs.gov/bod/18-01/#dmarc>

<https://cyber.dhs.gov/bod/18-01/#introduction-to-email-authentication>

<https://www.proofpoint.com/us/resources/white-papers/getting-started-with-dmarc>

Attack: We consider the attack to be the sending of a malicious email and refer to success or successes meaning a subset. (At least one quote seems to use attack to mean success: “became a victim of a BEC attack...”)

Attempt: Used synonymously with attack.

Business Email Compromise (BEC): Emails sent to trick a person into initiating a bank transfer. Often distinguished from other malicious emails by a lack of URL or attachment, making traditional defenses far less effective.

Display Name Deception: The use of an email display name to deceive. The display name is the name that most email clients will display, or the text that is not part of the email address. For example, if we look at “John Doe” <johndoe@example.com>, John Doe would be the display name. There are few restrictions on the contents of that string, it could be “alerts@fbi.com” or “Phil -- Global Cyber Alliance” <johndoe@example.com>. These are also called friendly name fraud and a variety of other names.

Domain-spoofed Email: An email that claims it comes from “bank.com” when it really comes from some other system not authorized to send bank.com email.

Email Address: A string of the form “person@example.com.” (This definition is incomplete, the string can also be <pedant@example.com> and other things which are not relevant here.)

Fraudulent Email

Human Action

Look-alike Domain: A domain designed to fool a person into thinking it relates to a different domain; the canonical example is paypa1.com, spelled with the number ‘1’ replacing the letter ‘l’. Also called cousin domain, or visually similar or homomorph domains. Advanced versions of this take advantage of unicode to display visually indistinguishable glyphs. See (Zheng) for an example.

Malicious Email

Secure Email Gateway (SEG): A type of product which “provides basic message transfer agent functions; inbound filtering of spam, phishing, malicious and marketing emails; and outbound data loss prevention (DLP) and email encryption.” (Gartner)

Spoofed Email: An email which purports to be from example.com, but was not sent from that domain’s authorized senders.

REFERENCES

(ABA) Corporate Account Takeover/Business Email Compromise, American Bankers Association, last visited June 23, 2018, <https://www.aba.com/Tools/Function/Fraud/Pages/CorporateAccountTakeover.aspx>

(Agari) Threat Taxonomy: A Working Framework to Describe Cyber Attacks, Agari, July 24, 2017, <https://www.slideshare.net/AgariData/threat-taxonomy-a-working-framework-to-describe-cyber-attacks>

(AGARI01) Business Email Compromise (BEC) Attack Trends Report, last visited Sept 16, 2018, <https://www.agari.com/resources/whitepapers/bec-report/>

(Amaroso) Fundamentals of computer security technology, Edward G. Amaroso (Prentice Hall, 1994)

(Anderson12) "Measuring the cost of cybercrime," Anderson, Ross, Chris Barton, Rainer Böhme, Richard Clayton, Michel JG Van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. In The economics of information security and privacy, pp. 265-300. Springer, Berlin, Heidelberg, 2013.

(Census) 2015 SUSB Annual Data Tables by Establishment Industry, United States Census Bureau, 2015 SUSB Annual Data Tables by Establishment Industry, US & states, totals, last visited June 28, 2018 <https://www.census.gov/data/tables/2015/econ/susb/2015-susb-annual.html>

(DMDatabases) USA BUSINESS LIST – EMPLOYEE SIZE PROFILE, Direct Marketing Databases.com, last visited June 24, 2018, <http://dmdatabases.com/databases/business-mailing-lists/how-many-businesses>

(FBI BEC) Business E-Mail Compromise, Federal Bureau of Investigation, February 27, 2017 <https://www.fbi.gov/news/stories/business-e-mail-compromise-on-the-rise>

(FBI17) 2017 Internet Crime Report, Retrieved June 26, 2018 from https://pdf.ic3.gov/2017_IC3Report.pdf

(FBI18) BUSINESS E-MAIL COMPROMISE THE 12 BILLION DOLLAR SCAM Federal Bureau of Investigation Alert # I-071218-PSA, July 12, 2018, <https://www.ic3.gov/media/2018/180712.aspx> last visited August 23, 2018

(GREATHORN) Spear Phishing Report, January 2017, retrieved from <https://cdn2.hubspot.net/hubfs/851792/Content%20for%20Resources%20Page/GreatHorn%20Spear%20Phishing%20Report%20-%202017.pdf>

(Home Office) Understanding the costs of cyber crime: A report of key findings from the Costs of Cyber Crime Working Group, Research Report 96, (UK) Home Office Science Advisory Council, January 2018

(IC3 BEC) BUSINESS E-MAIL COMPROMISE E-MAIL ACCOUNT COMPROMISE

THE 5 BILLION DOLLAR SCAM, Federal Bureau of Investigation, Alert # I-050417-PSA, May 4, 2017, <https://www.ic3.gov/media/2017/170504.aspx>

(Illumenmedia) 40 Unbelievable Email Marketing Stats Prove It's Thriving Into 2018, Brian Robben, November 16, 2017, <https://illumenmedia.com/email-marketing-stats/>

(Romanoski) Examining the costs and causes of cyber incidents, Sasha Romanoski, Journal of Cybersecurity, Volume 2, Issue 2, 1 December 2016, Pages 121-135, <https://doi.org/10.1093/cybsec/tyw001>

REFERENCES

(NetDiligence) 2017, 2016 Cyber Claims Study, NetDiligence, retrieved June 26, 2019 <https://netdiligence.com/portfolio/cyber-claims-study/>

(PROOF01) “85% of Organizations targeted by at least one business email compromise (BEC) attack in Q1 2017”, retrieved from <https://www.proofpoint.com/us/corporate-blog/post/85-organizations-targeted-least-one-business-email-compromise-bec-attack-q1-2017>

(PROOF17) Email Fraud Threat Report, Year in Review. Retrieved from <https://www.proofpoint.com/sites/default/files/pfpt-us-tr-email-fraud-yir-180212.pdf>

(Riek) “Estimating the costs of consumer-facing cybercrime: A tailored instrument and representative data for six EU countries (working paper)” Riek, Markus, Rainer Boehme, Michael Ciere, Carlos Gañán, and Michel van Eeten. In Workshop on the Economics of Information Security (WEIS). 2016.

(Snow) Gordon Snow, “Cyber Security: Threats To The Financial Sector”, Congressional Testimony, September 14, 2011, <https://financialservices.house.gov/uploadedfiles/091411snow.pdf>

(Statista) E-mail marketing spending in the United States from 2014 to 2019 (in billion U.S. dollars), last visited June 26, 2018, <https://www.statista.com/statistics/266624/e-mail-marketing-expenditure-in-the-united-states/>

(SYMC) Email Threats 2017: An ISTR Special Report. October, 2017. Ben Nahorney. <https://www.symantec.com/content/dam/symantec/docs/security-center/white-papers/istr-email-threats-2017-en.pdf>

(TrendMicro) TrendLabs 2016 Security Roundup: A Record Year for Enterprise Threats. <https://documents.trendmicro.com/assets/rpt-2016-annual-security-roundup-a-record-year-for-enterprise-threats.pdf>

(Valimail)

Dylan Tweney, “Man Bites Dog? Federal Government Leads in DMARC Adoption”, blog post, last visited Aug 23, 2018, <https://www.valimail.com/blog/email-fraud-q1-2018/>

(ValimailFraud) The Fake Email Crisis: 6.4 Billion Fake Messages Every Day, Dylan Tweney, 2018 <https://www.valimail.com/resources/report/email-fraud-landscape-q2-2018/>

(Verizon15) 2015 Data Breach Investigations Report, https://www.verizonenterprise.com/resources/reports/rp_data-breach-investigation-report_2015_en_xg.pdf

(Verizon16) 2016 Data Breach Investigations Report, https://www.verizonenterprise.com/resources/reports/rp_DBIR_2016_Report_en_xg.pdf

(Verizon17) 2017 Data Breach Investigations Report, https://www.verizonenterprise.com/resources/reports/2017_dbir_en_xg.pdf

(Verizon18) 2018 Data Breach Investigations Report, https://www.verizonenterprise.com/resources/reports/rp_DBIR_2018_Report_en_xg.pdf

(Wombat/Proofpoint 2018)

Wombat/Proofpoint State of the Phish Report 2018, <https://www.wombatsecurity.com/state-of-the-phish>

(Zheng) Phishing with Unicode Domains, Xudong Zheng, April 14, 2017 <https://www.xudongz.com/blog/2017/idn-phishing/>

About The Authors

Shostack & Associates is a specialized security consultancy, focused on meeting the unique needs of each client through a variety of services including threat modeling, security engineering and risk management.

Adam Shostack

is a consultant, entrepreneur, technologist, author and game designer. He's a member of the BlackHat Review Board, and helped found the CVE and many other things. He's currently helping a variety of organizations improve their security, and advising startups as a Mach37 Star Mentor. While at Microsoft, he drove the Autorun fix into Windows Update, was the lead designer of the SDL Threat Modeling Tool v3 and created the "Elevation of Privilege" game. Adam is the author of Threat Modeling: Designing for Security, and the co-author of The New School of Information Security.

The Cyentia Institute

is a Virginia-based research services firm that exists to advance cybersecurity knowledge and practice through use-inspired, data-driven research. We curate and publish research for the community, partner with other organizations to create compelling publications and help enterprises turn complex security data into confident strategic decisions.

Jay Jacobs

is a Co-Founder of and Chief Data Scientist at Cyentia Institute, a research firm dedicated to advancing the state of information security knowledge and practice through data-driven research. He is best known for contributions to Verizon's annual Data Breach Investigations Report series and his book "Data-Driven Security: Analysis, Visualization and Dashboards." He is a founding member of the Society of Information Risk Analysts, and remains an active proponent of improving how we measure and manage risk.

Wade Baker

is a Co-Founder of the Cyentia Institute, which focuses on improving cybersecurity knowledge and practice through data-driven research. He's also a professor in Virginia Tech's College of Business, teaching courses for the MBA and MS of IT programs. Prior to this, Wade held positions as the VP of Strategy at ThreatConnect and was the CTO of Security Solutions at Verizon, where he had the great privilege of leading Verizon's annual Data Breach Investigations Report (DBIR) for 8 years.

About the Global Cyber Alliance

The Global Cyber Alliance (GCA) is an international, cross-sector effort dedicated to eradicating cyber risk and improving our connected world. We achieve our mission by uniting global communities, implementing concrete solutions, and measuring the effect. GCA, a 501(c)3, was founded in September 2015 by the Manhattan District Attorney's Office, the City of London Police and the Center for Internet Security.

Learn more at www.globalcyberalliance.org.

