ACID | AGARI CYBER INTELLIGENCE DIVISION

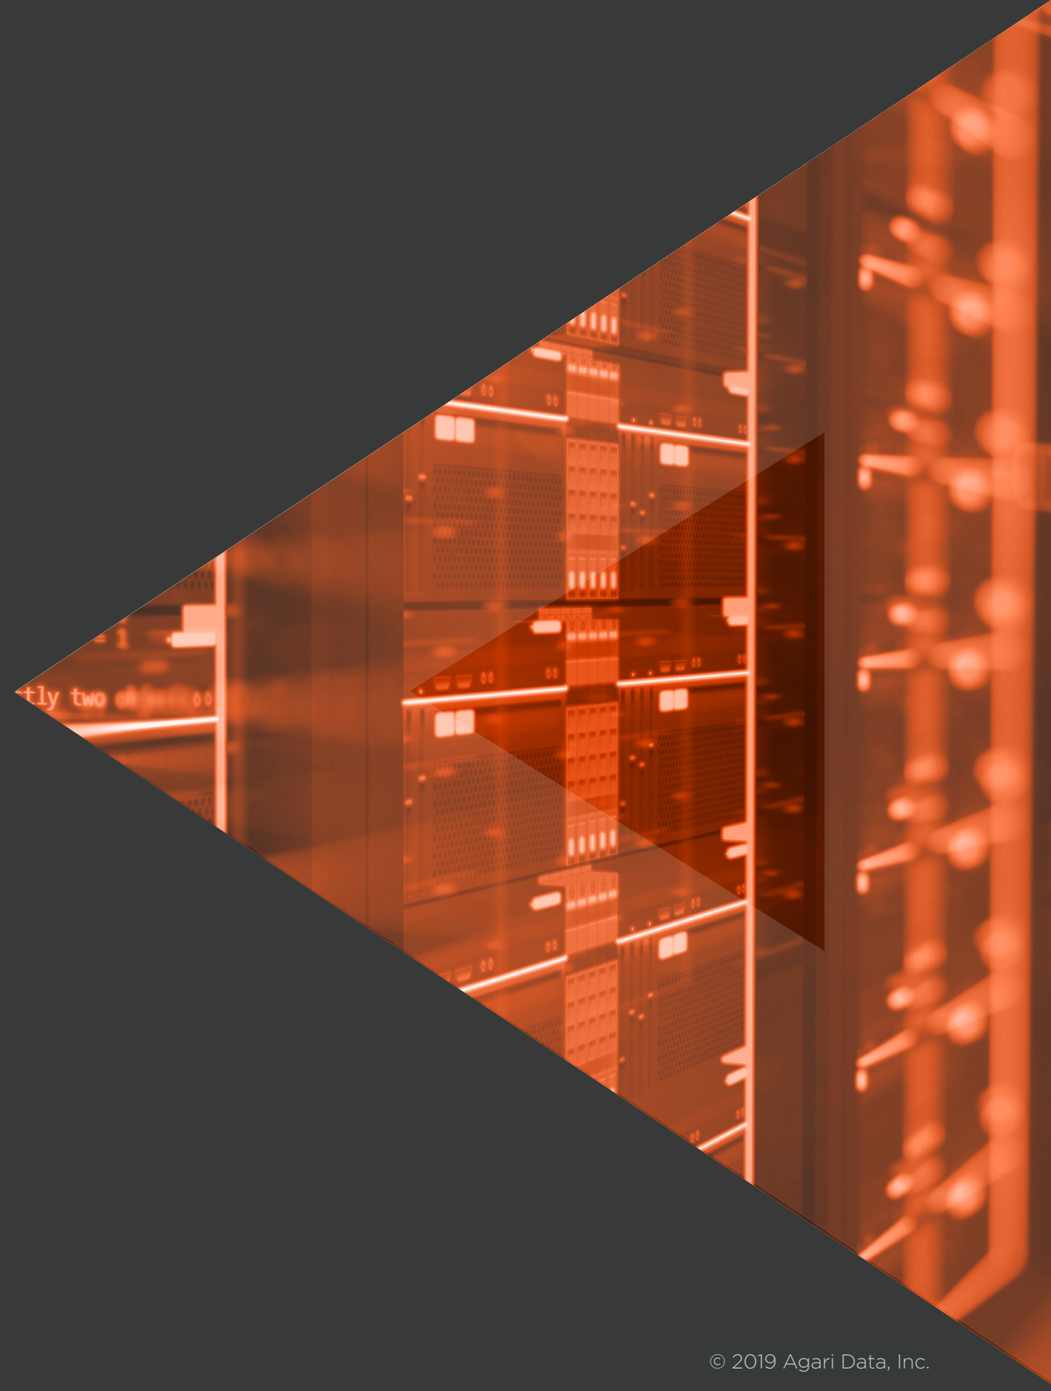

REPORT

Q1 2019

Email Fraud & Identity Deception Trends

Global Insights from the Agari Identity Graph™

AGARI

Email-Based Fraud and Identity Deception are Evolving Fast

Email remains the killer app for communication and collaboration in both business and everyday life. But it's under attack like never before. A lack of built-in authentication has long given fraudsters the ability to send an email claiming to be someone else. But today, a new generation of cybercriminal organizations is the driving force behind rapidly-evolving, socially-engineered email threats that grow more dangerous by the day.

Evil in the Inbox

Over the past year, business email compromise (BEC) scams have jumped 60%. More than 90% of organizations report being hit by targeted email attacks, with 23% suffering financial damage that can average $1.6 million and up. 96% of successful data breaches now begin with an email, wreaking an average $7.9 million in costs per incident.

What is driving this uptick? Increasingly sophisticated cybercriminal organizations that pair identity deception techniques with personalized, socially-engineered emails designed to throw recipients off-kilter just long enough to fork over login credentials or make wire transfers before thinking to confirm the message's legitimacy. Despite increased awareness of the problem, the price tag is estimated at $12.5 billion—and counting.[1]

Hijacking Your Brand, Targeting Your Consumers

Businesses aren't alone in the crosshairs. Every minute of the day, 22.9 new phishing attacks target consumers by impersonating trusted brands. Whether it's through a fake "payment past due" or a "fraud alert" email, these and other Internet scams bamboozle consumers out of $1.4 billion through brand impersonation each year.

KEY FINDINGS

- Account takeover-based threats account for 20% of the inbound attacks that target employees.
- While 70% of brand impersonation attacks spoofed Microsoft, another notable impersonation target was the IRS.
- Costs reported to Security Operations Centers (SOCs) exceeded $4.86M to triage, investigate, and remediate.
- The volume of raw DMARC domains surged to 6.1 million, but major businesses are still lagging in adoption rates.

[1] Sources: Agari, Verizon, Ponemon, FBI.

Inside this Report

In this report, we look at trends in phishing and email fraud against business, as well as those targeting their customers through domain spoofing and other tactics. For the first time, we examine the impact of phishing incident response by tracking the burden and cost for a SOC team to respond to user-reported emails. The statistics presented here reflect information captured from the following sources over the fourth quarter—October through December—of 2018:

- Data extracted from the 300 million+ daily model updates by the Agari Identity Graph
- DMARC-carrying domains identified within the 330 million+ domains crawled
- Insights captured from a phishing incident response survey of over 300 cybersecurity professionals

The Agari Cyber Intelligence Division (ACID) is the only counterintelligence research team dedicated to worldwide BEC and spear phishing investigation. ACID supports Agari's mission of protecting communications so that humanity prevails over evil. The ACID team uncovers identity deception tactics, criminal group dynamics, and relevant trends in advanced email threats. Created by Agari in 2018, ACID helps to impact the cyber threat ecosystem and mitigate cybercrime activity by working with law enforcement and other trusted partners.

Table of Contents

Employee Phishing and Business Email Compromise

Speak of the Devil:
A Taxonomy of Advanced Email Threats

With increasing levels of cybercrime posing a serious threat to individuals, businesses, and governments, it's vitally important to codify a consistent set of terms to describe the different challenges that make up this threat. Not every email scam is a "phishing attack," for instance.

To address this need, ACID has established a classification system for cyber threats—a threat taxonomy—that breaks down common email-based attacks in terms of how they are carried out, and what the perpetrators aim to achieve. This taxonomy will help readers understand the terms used in this report and what they mean to email security.

Because email fraud centers around identity deception, or the impersonation of trusted senders in order to con recipients, we start with the method by which the impostor impersonates the trusted sender's email account—making it appear as if the emails the impostor is sending are originating from the trusted party.

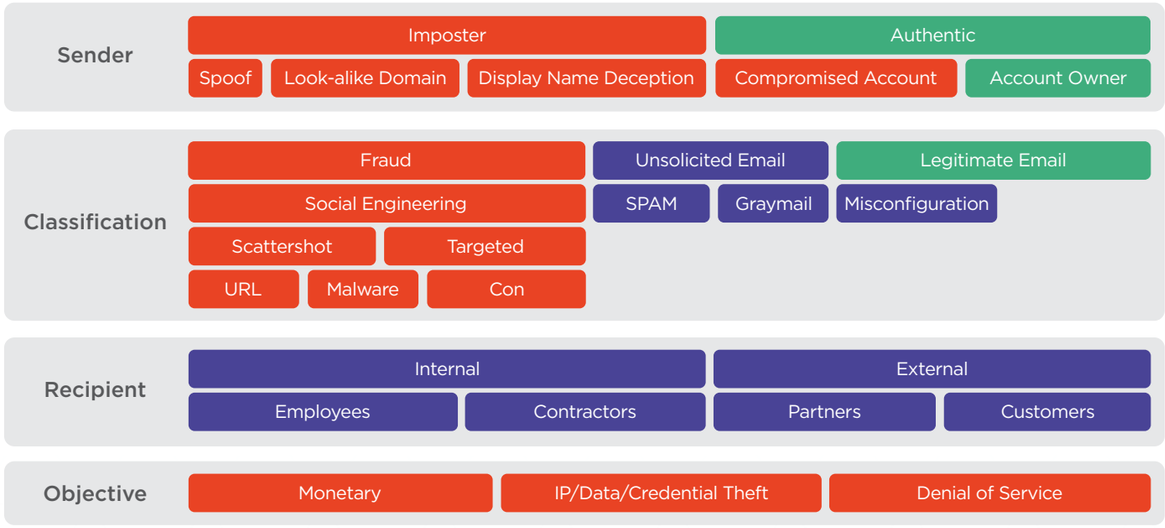

For more information about the Agari Threat Taxonomy, see **agari.com/taxonomy**

KEY FINDINGS

- Compromised accounts were used in 20% of identity deception attacks, showcasing the need for tighter security.
- One-third of attacks targeting C-level executives employ display name deception impersonating specific individuals.
- Microsoft and its related services continue to lead the way in brand impersonation attacks.

Leading Attack Modalities

Generally speaking, we observed three primary ways in which cybercriminals impersonate an email account:

LOOK-ALIKE DOMAINS AND DOMAIN SPOOFING: With look-alike domains, the cybercriminal registers a domain that is very similar to the legitimate domain he or she is seeking to impersonate. Look-alike domains are distinguished from domain spoofing, in which the attacker uses the actual email address of the impersonated identity in the "From" header—for example, "Company Customer Service" <noreply@company.com>. Email authentication standards such as DMARC can be used by a domain owner to prevent spoofing of the domain, but are still not adopted widely by all businesses. Domain spoofing is addressed in Part 3 of this report.

DISPLAY NAME DECEPTION: The cybercriminal inserts the name of the impersonated individual or brand into the "From" field within Gmail, Yahoo, or another free cloud-based email platform. These are also known as "friendly from" attacks.

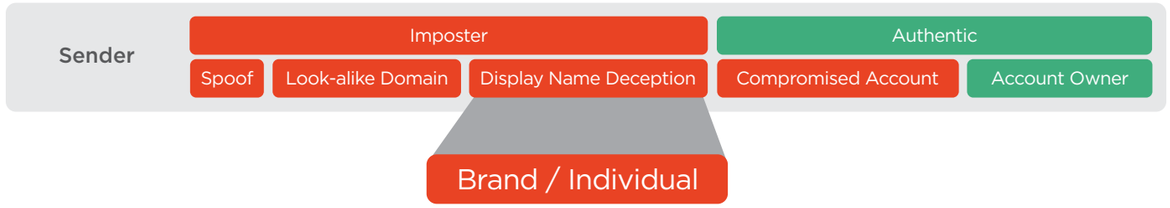

COMPROMISED ACCOUNT ATTACKS: The cybercriminal sends targeted requests from an account that's already been compromised—assuming the identity and the actual email account of the impersonated individual or brand, which is the most dangerous threat of all.

Different types or classes of attacks will entail different elements of this taxonomy.

A business email compromise (BEC) attack, for instance, can involve an impostor who aims to impersonate a trusted individual or brand using a look-alike domain, display name deception, or in the worst cases, a compromised legitimate account, leveraging sophisticated social engineering tactics to send highly personalized attacks. Impersonated individuals may be executives within the target's own company, or an outside vendor or partner company. A BEC attack is targeted and uses a con with no URL or attachment.

By comparison, a phishing attack may use any identity deception technique and send more broad-based messages meant to fool someone into clicking on a malicious link that captures their username and password. When attacking businesses, display name deception is the tactic of choice for cybercriminals seeking to impersonate the email account of a trusted individual or brand.

Patterns of Deceit:
Compromised Accounts Account for 20% of Attacks

Attacks launched from the hijacked email accounts of trusted individuals and brands hold the potential for major offensives in the months ahead.

Fake ‘From’ Lines: Shift Happens

Display name deception continues to be the tactic of choice for cybercriminals, accounting for 63% of all identity deception-based email attacks aimed at impersonating a trusted individual or brand—typically an outside vendor, supplier, or partner. But as this approach continues to gain traction over look-alike domains and simple domain spoofing, the nature of these impersonations appears to be in transition.

Fraudsters continue to favor impersonating trusted brands (50%) over trusted individuals (13%). But it’s notable that this report reflects a slight drop in brand impersonations from the previous quarter. It also corresponds with a 61% jump in impersonations of trusted individuals, up from just 8% in 90 days.

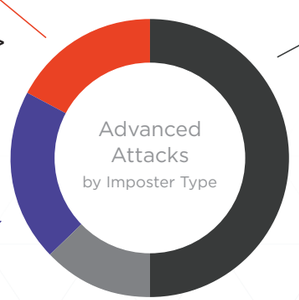

17%
Look-alike Domain

From: LinkedIn <noreply@liinkediin.com>
To: Jan Bird <jan.bird@gs.com>
Subject: Diana has endorsed you!

50%
Display Name Deception (Brand)

From: Chase Support <chase@gmail.com>
To: Tom Frost <ffrost@amazon.com>
Subject: Account Disabled

20%
Compromised Account

From: Raymond Lim <rlim@contoso.com>
To: Cong Ho <cho@contoso.com>
Subject: PO 382313

13%
Display Name Deception (Individual)

From: Ravi Khatod <Ravi Khatod [hackyjoe@gmail.com]>
To: Cong Ho <cong@agari.com>
Subject: Follow up on Invoice Payment

Notably, compromised accounts were used in 20% of identity-deception attacks. Legitimate email accounts that have been taken over by scammers can be an effective method to distribute phishing emails because they are, in a sense, trusted—allowing them to bypass mail filters more easily.

A potential driver for the proportion of attacks attributed to ATO-based email attacks could be the continually expanding marketplace on the dark web for stolen login credentials belonging to high-value targets.

The impact of this attack type cannot be overstated. Attacks launched from compromised email accounts are by far the hardest to detect and disrupt, making them a serious vulnerability for the account's legitimate owner and the companies involved.

Indeed, a successful account takeover does not just give fraudsters the ability to impersonate the account's owner. It also gives them access to the individual's contacts, ongoing email conversations, and historical email archives—making it possible to craft new scams made all the more galling by their extraordinary personalization and crushing effectiveness.

The remaining 17% of identity-deception emails use look-alike domains to send malicious content. While some of these domains can be simply spoofed and sent from basic mailing tools, others are registered by phishing threat actors. The cost associated with registering a domain reduces a scammer's overall return on investment, which is why this tactic likely is not used more frequently. Why pay for infrastructure when you can create a free, temporary email account—especially when the success rate is likely the same?

Plenty of Phish in the C-Suite:
Display Name Deception a Key Tactic When Impersonating Executives

Identity deception trendlines take on vastly different trajectories when filtered for attacks targeting senior executives. In the fourth quarter, for instance, display name deception impersonating a specific individual constituted only 13% of attacks targeting the general employee population.

But fully one-third of attacks targeting C-level executives employ this tactic. This variance is driven by the volume of BEC scams targeting CFOs and other senior finance executives.

Given that the objective of these schemes is to manipulate recipients into initiating wire transfers, malicious email messages appearing to come from the CEO and other C-suite executives can inspire prompt action—indicating this may be one of the primary email threats facing senior executives.

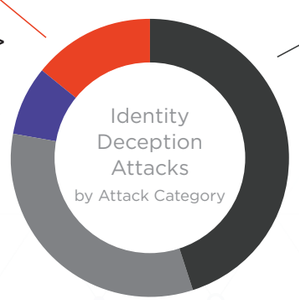

14%
Look-alike Domain

From: LinkedIn <noreply@liinkediin.com>
To: Jan Bird <jan.bird@gs.com>
Subject: Diana has endorsed you!

45%
Display Name Deception (Brand)

From: Chase Support <chase@gmail.com>
To: Tom Frost <ffrost@amazon.com>
Subject: Account Disabled

8%
Compromised Account

From: Raymond Lim <rlim@contoso.com>
To: Cong Ho <cho@contoso.com>
Subject: PO 382313

33%
Display Name Deception (Individual)

From: Ravi Khatod <Ravi Khatod [hackyjoe@gmail.com]>
To: Cong Ho <cong@agari.com>
Subject: Follow up on Invoice Payment

Meanwhile, compromised email accounts are leveraged only sparingly for attacks targeting senior executives, accounting for only 8% of attacks during the last three months of 2018. More targeted research and personalization may account for the fact that executives seem to be far more lucrative to fraudsters if the high-value target accounts can be compromised and used to launch attacks targeting employees that rank lower on the organizational chart.

For more information on how cybercriminals target the C-level, see **agari.com/londonblue**

More Fraudsters Masquerade as Microsoft:
#1 Most Impersonated Brand Grows More Popular

Microsoft and its business units remain cybercriminals' go-to disguise when impersonating brands.

Microsoft services continue to lead the way in brand impersonation attacks, consistent with trends seen over the last few years. During the final three months of 2018, 44% of brand deception attacks displayed a Microsoft service as a way to deceive victims—up from 36% in the third quarter of 2018. As the chart below indicates, the last quarter of 2018 featured a strong showing by the IRS from an impostor perspective.

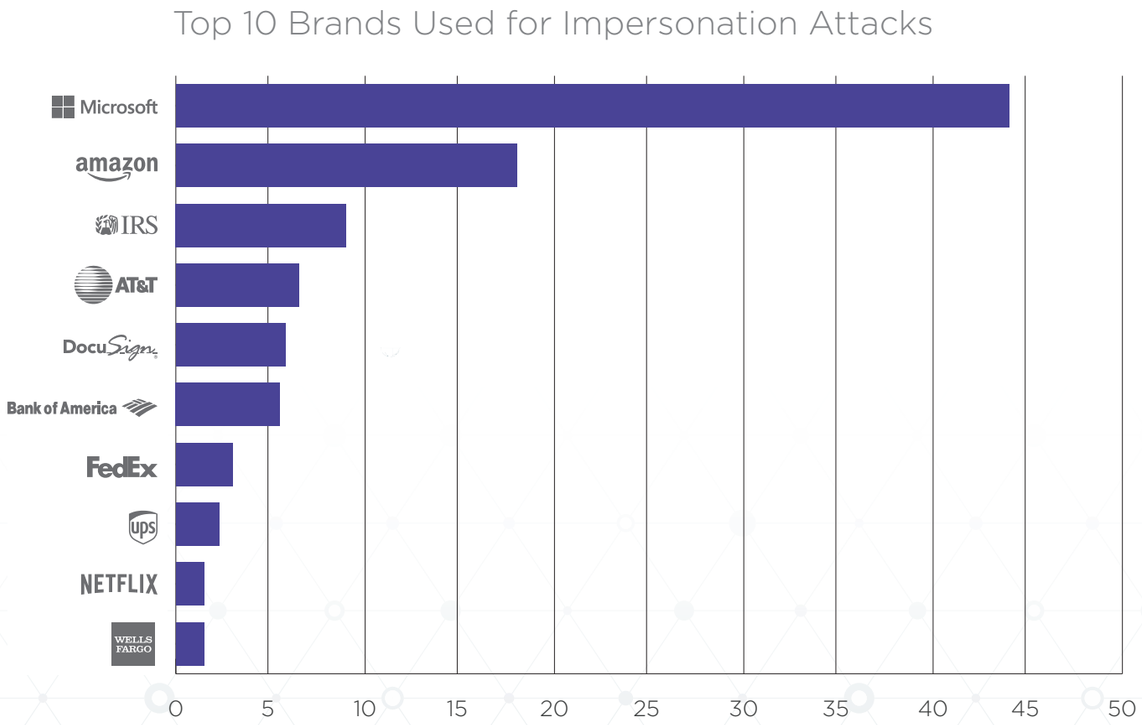

As 2018 drew to a close, deception attacks impersonating the Internal Revenue Service (IRS) shot upward. Driven by the annual scourge of BEC scams aimed at stealing W-2 information in the run-up to tax filing season, nearly one in ten identity deception-based emails impersonated the IRS, up from two percent in the third quarter.

For attacks targeting high-value executive targets, Microsoft remains the top target—accounting for more than 8 out of 10 brand impersonations. Trailing far behind in second place is FedEx, followed by the IRS and UPS. Shipping services are a common impersonation target, especially around the holiday season, because the delivery of packages during this time of year is expected,—making the phishing emails more contextually appropriate.

In most attacks involving these or other most impersonated brands, phishing scams are likely aimed at credentials harvesting, in hopes of hijacking accounts from which to launch more highly-targeted attacks of all kinds, including wire fraud-based BEC schemes.

Leveraging the vulnerability Secure Email Gateways have in detecting and mitigating display name deception, attackers attempt to exploit the relationship employees have with trusted technology and financial brands.

A case in point was the carefully crafted Amazon brand impersonation sent to an AWS admin at a software/SaaS company. This credential phishing attack was especially pernicious given the dependence many enterprises place on web and compute services provided by AWS.

Brand Impersonation Attack Example: AWS Credential Phishing

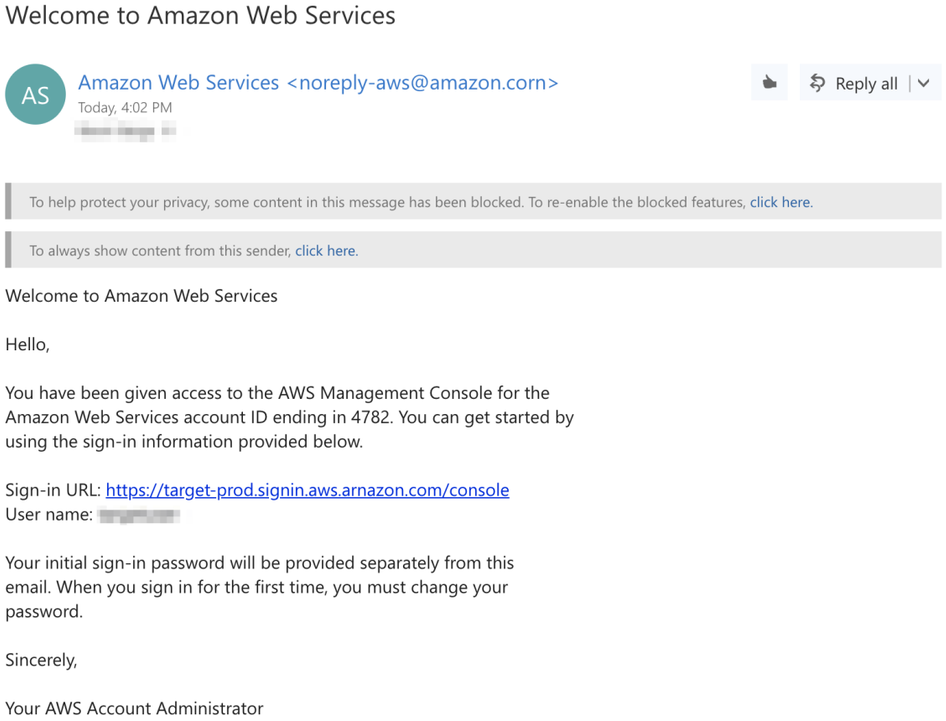

Welcome to Amazon Web Services

AS Amazon Web Services <noreply-aws@amazon.corn>
Today, 4:02 PM
[redacted]

Reply all

To help protect your privacy, some content in this message has been blocked. To re-enable the blocked features, click here.

To always show content from this sender, click here.

Welcome to Amazon Web Services

Hello,

You have been given access to the AWS Management Console for the Amazon Web Services account ID ending in 4782. You can get started by using the sign-in information provided below.

Sign-in URL: https://target-prod.signin.aws.arnazon.com/console
User name: [redacted]

Your initial sign-in password will be provided separately from this email. When you sign in for the first time, you must change your password.

Sincerely,

Your AWS Account Administrator

Phishing Incident Response Trends

Employee-Reported Phishing is Flooding Security Operations Centers

While businesses strive to implement security controls to prevent phishing emails from reaching employee inboxes, there will always be a risk that employees will receive malicious emails intended to defraud the company or steal sensitive information as part of a data breach. For US-based companies, the average cost of a breach now runs $7.9 million, and the probability of a breach occurring is now 14% per year, according to the Ponemon Institute Cost of Data Breach Study 2018.

Employee Reporting as Threat Intelligence

With the vast majority of businesses implementing security awareness training, phishing simulation, and the ability for employees to report phishing, it's critical to understand how to leverage this threat feed to discover and contain breaches before data is exfiltrated. To that end, it's crucial for businesses to streamline the process of triaging, investigating, and remediating phishing incidents to avoid flooding the security operations center (SOC) with more phishing incidents to investigate than it can handle. Otherwise, intelligence regarding breaches may go undiscovered until it is far too late.

Phishing Incident Response Survey

As part of the introduction of Agari Incident Response to the market, ACID conducted a survey of 325 organizations ranging in size from 1,000 employees to 209,000 employees. Of the respondents, 237 were based in the United States with 83 based in the United Kingdom.

The respondents included a combination of both Agari customers and non-customers—74 and 251 respectively. The survey asked a series of questions regarding employee-reported phishing—including reporting mechanism, volume, false positive rate, existing tools for phishing incident response, and time required to investigate phishing. This section of the Q1 2019 report highlights analysis of responses to these questions.

KEY FINDINGS

- Employees reported an average of 23,063 phishing incidents to the Security Operations Center each year.
- SOC analysts spent an average of 3.96 hours on a false positive, and 5.88 hours on a valid phish.
- Costs reported to the Security Operations Centers exceeded $4.86M to triage, investigate, and remediate.

Empowering Employees:
Impact of Security Awareness Training on Phishing Reporting

Respondents report that 98% of employees have the ability to report phishing attacks, and often even have a convenient button and/or abuse inbox to forward suspicious messages to the security team. Eighty-eight percent of organizations report using a phishing simulation vendor to test employees' ability to detect a phishing incident after participating in security awareness training.

Training Employees to Report Phishing

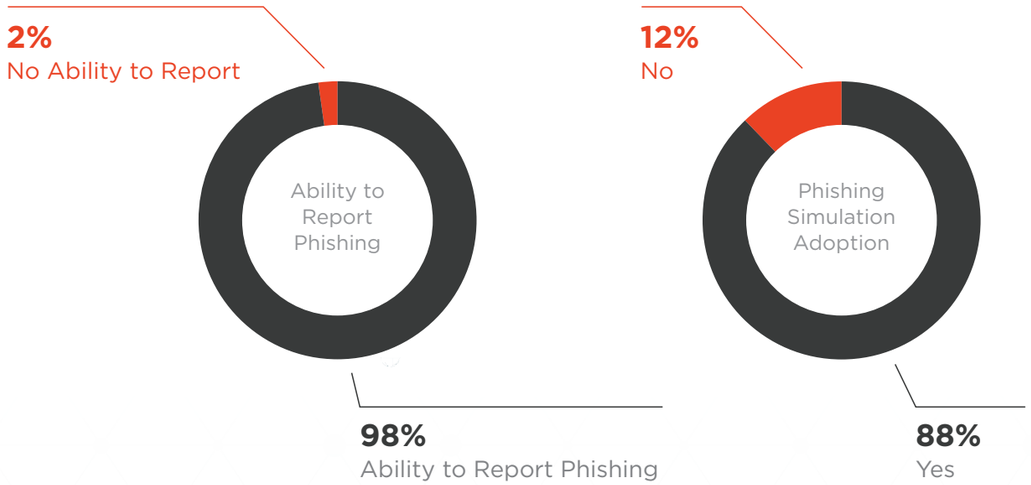

Hitting the Panic Button:
How Do Employees Report Phishing?

While the most common method available to employees to report phishing is an abuse@company.com inbox, most companies offer multiple other methods, including filing a help desk trouble ticket, using the native email client such as the Microsoft Office 365 example, or using a third-party email client button like the KnowBe4 phishing button example.

Employee Options to Report Phishing (Global)

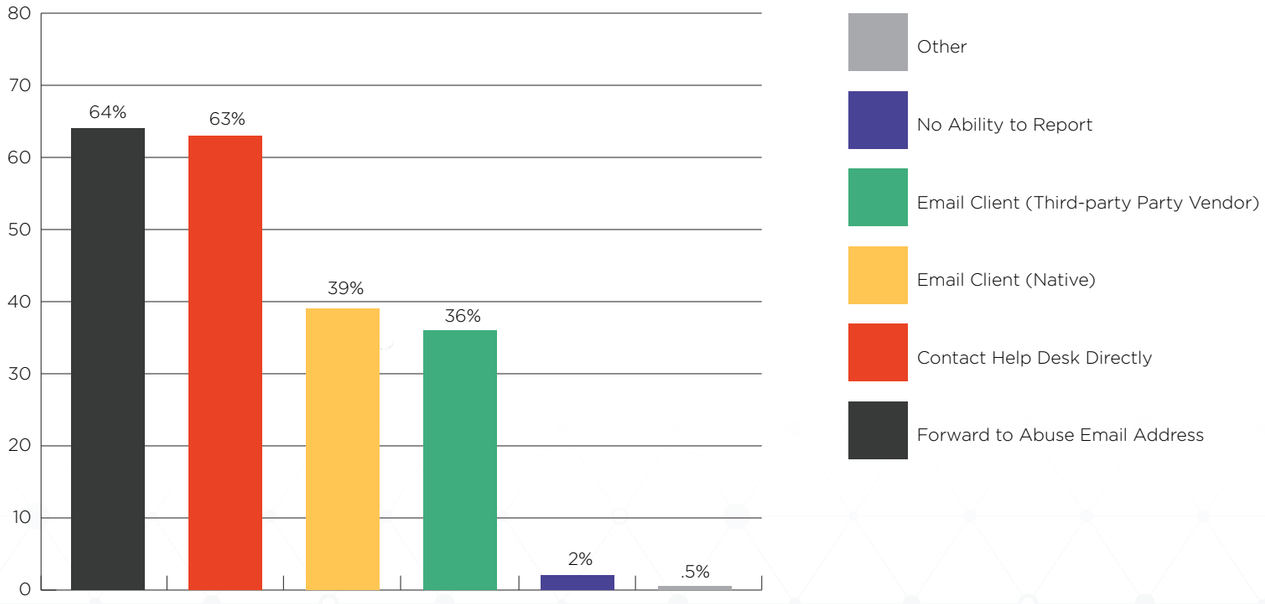

Whether the phishing incident is reported through an inbox or phishing button, the phishing email itself is forwarded to some combination of a security operations center (SOC), help desk support center, for investigation and remediation team.

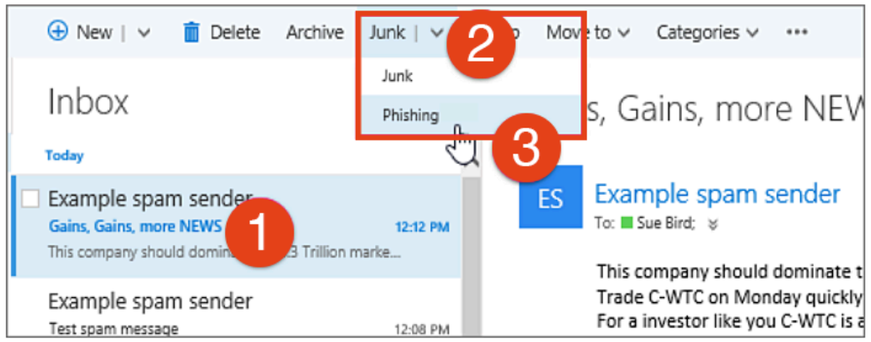

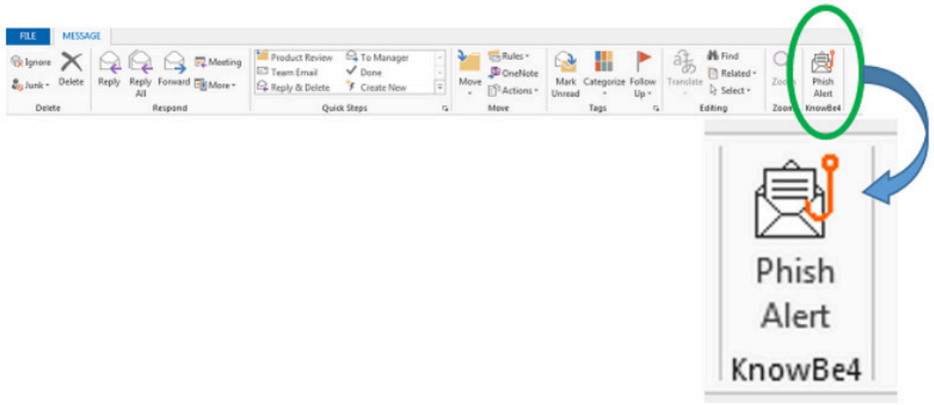

In some cases, the mail platform (Microsoft Office 365 or Google Suite) or phishing simulation vendor also receives a copy of the reported phishing messages.

Employee Reported Phishing Incidents Volume and Accuracy

With 98% of employees having the ability to report phishing and 88% being tested regularly on their ability to identify phishing incidents, the next logical question to answer is “What is the volume of employee reported phishing incidents?”

Based on the 308 organizations we surveyed—222 in the United States and 82 in the United Kingdom—employees report more than 23,000 phishing incidents per organization on an annual basis, with a slightly higher number of phishing incidents in UK-based companies.

Volume Per Organization of Phishing Incidents

Average Number of Reported Phishing Incidents Per Organization Annually

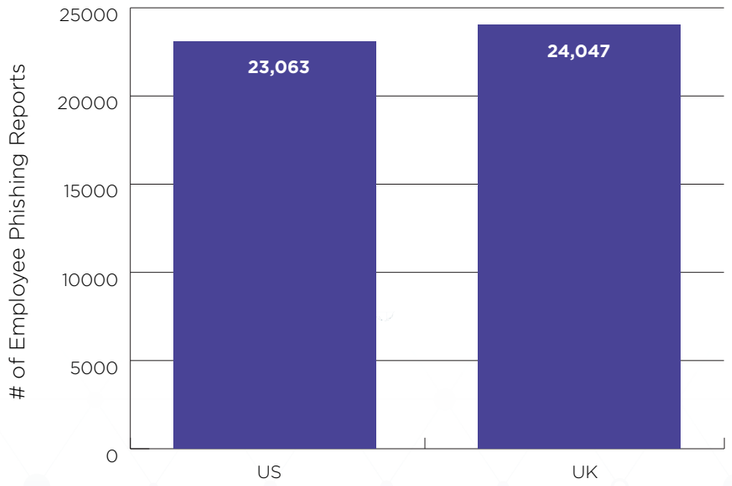

Distribution of Annual Reported Phishing Incidents (Global)

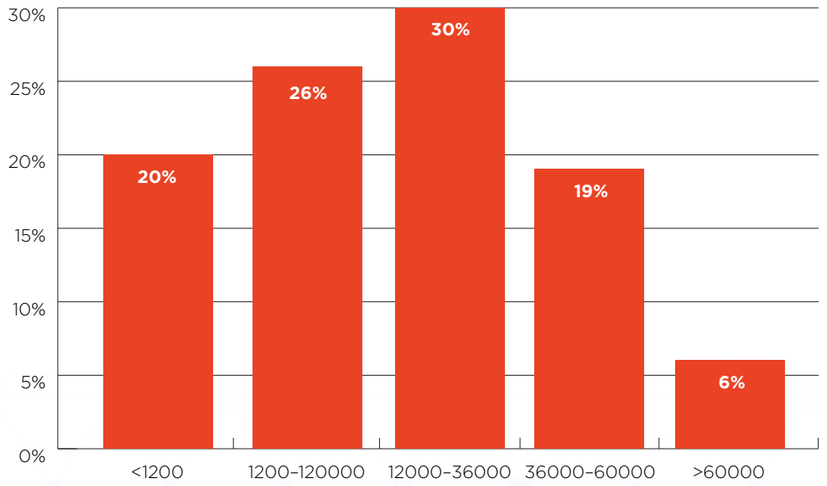

30% of respondents reported phishing incidents to be between a common range of 12,000 to 36,000 per year.

False Positive Rate

While employees frequently report phishing, the emails they report are not always true phishing incidents.

Security training often encourages users to report any suspicious email. As a result, spam, unwanted marketing emails, as well as legitimate email is often reported as phishing—even when they are not.

When we asked organizations “what percentage of employee phishing reports were determined to be false positives?” companies reported that their false positive rate was 50% on average, with a slightly higher false positive rate in UK-based companies.

Employee Reported Phishing False Positive Rate

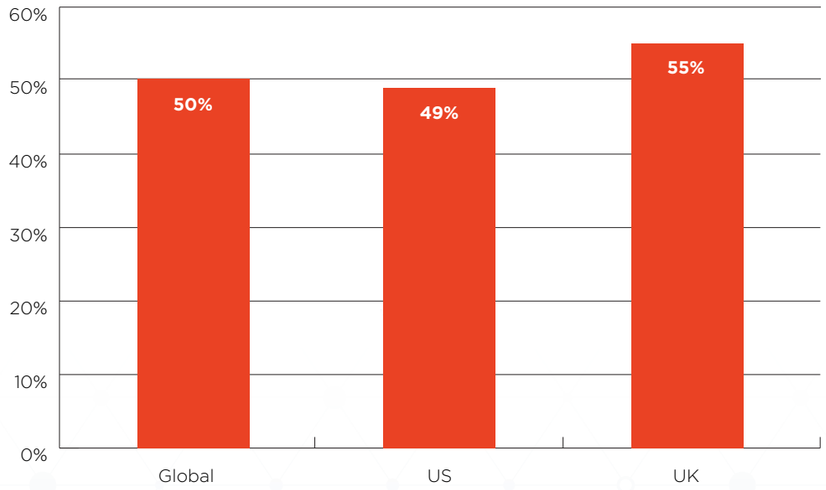

Time Required for Triage, Investigation, Forensics, and Remediation

Phish Reporting

Employees report suspect message using phish button

SOC Triage

SOC handles reports, filtering out obvious false positives

Forensic Analysis

SOC Analyst determines level of impact

Incident Remediation

SOC works with Messaging to address incidents

In the survey, respondents were asked: “For employee phishing reports, how much time on average does it take a SOC analyst to triage, investigate, and remediate?” This question was asked in the context of both true phishing incidents and false positive reports.

The overall average across all phishing incidents was 4.9 hours to triage, investigate, and remediate. On average, SOC analysts spent 3.96 hours triaging a false positive, and 5.88 hours triaging, investigating, and remediating a valid phish.

The triage process typically involves a quick investigation of the sender domain and address, URLs, and attachment to determine if the message is potentially malicious. This process is often manual, requires multiple third-party tools, and involves the judgement of the analyst.

By comparing the average false positive to true phish time, we estimate that 67% of SOC analyst time is spent in the triage phase of the process, while only 33% is spent on forensic analysis and remediation.

Average Time Per Phishing Incident to Triage. Investigate, and Remediate

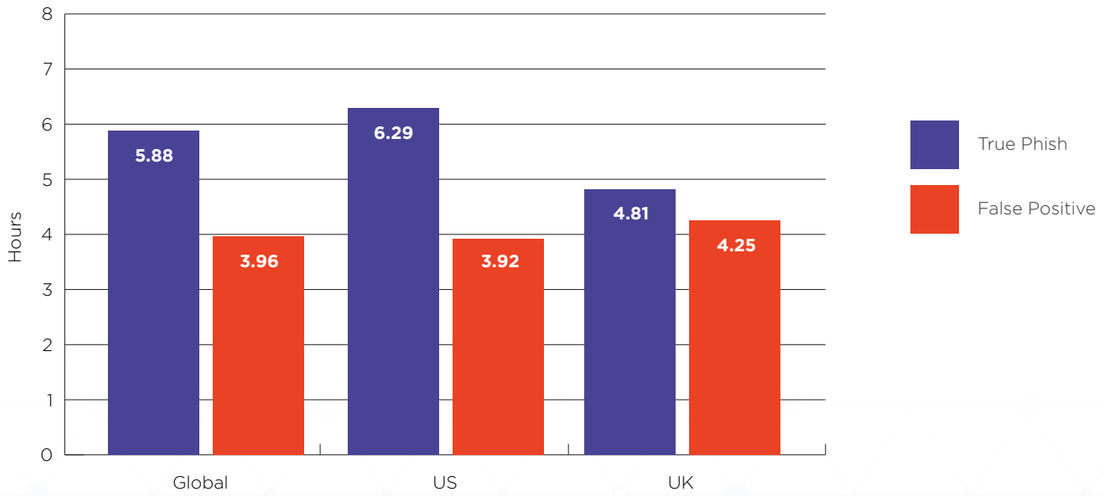

Reality Check:
A SOC Staffing Snapshot

To determine if SOCs are adequately staffed to handle phishing incidents in a timely manner, respondents were asked about the size of the SOC team.

A full 94% of organizations reported having at least one dedicated SOC analyst. As you might expect, the analysis showed a strong correlation between company size, the number of phishing incidents, and the number of SOC employees.

For example, 41% of organizations with more than 10,000 employees had 20 or more SOC analysts. The same is true of organizations with 60,000 or more phishing incidents per year.

The Staffing Gap

Based on the average number of phishing incidents and the average time to remediation (4.9 hours), the average SOC needs 54 analysts to handle the number of phishing incidents per company. Given that the average number of SOC analysts in our survey is 12.5, there is a staffing gap of at least 41.5 full-time equivalents (FTEs). This gap currently results in most organizations failing to detect phishing incidents, which opens each organization to the possibility of breaches or fraud.

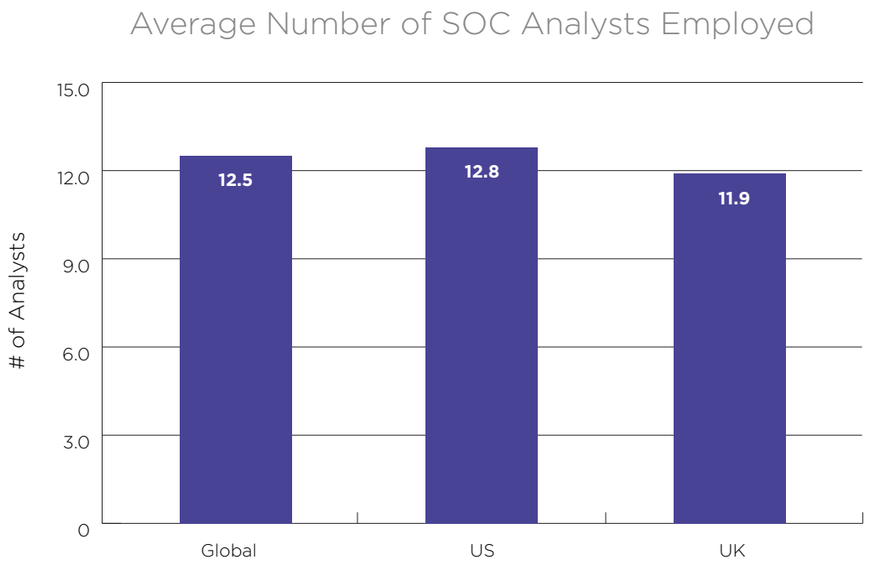

Rising Costs:
Data Breach Economics

According to the 2018 Verizon Data Breach Investigations Report (DBIR), the entry point for 96% of data breaches is email. The average cost of a data breach in the United States is $7.9 million, with a 14% probability of a breach occurring annually, according to Ponemon Institute. If you multiply the average breach cost of $7.9 million by the probability of 14%, the annual breach risk is $1.1 million.

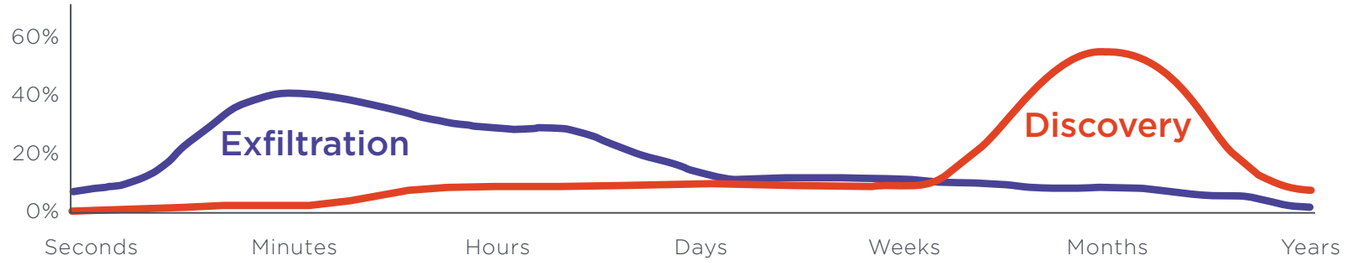

Source: 2018 Verizon DBIR

Meanwhile, the DBIR finds that the average data breach results in exfiltration of data within minutes or hours—while the average time-to-discovery takes months. This is likely a symptom of understaffed and inefficient SOC processes for handling phishing incidents. Ideally, SOC analysts would be able to triage, investigate, and remediate reported phishing incidents within minutes, enabling the business to remediate the compromise and contain the breach.

Automate or Else?
Mitigating Breach Risk by Reducing Time-to-Remediation

As part of the phishing incident response survey, we asked respondents how much reducing the response time required for phishing incident response would reduce their breach risk. Overall, businesses felt they could reduce breach risk by 50% by automating the process of phishing incident response.

A 50% reduction in breach risk would result in a $551,025 decrease in annual breach risk for the average business.

Risk Reduction Due to Automated Phishing Incident Response

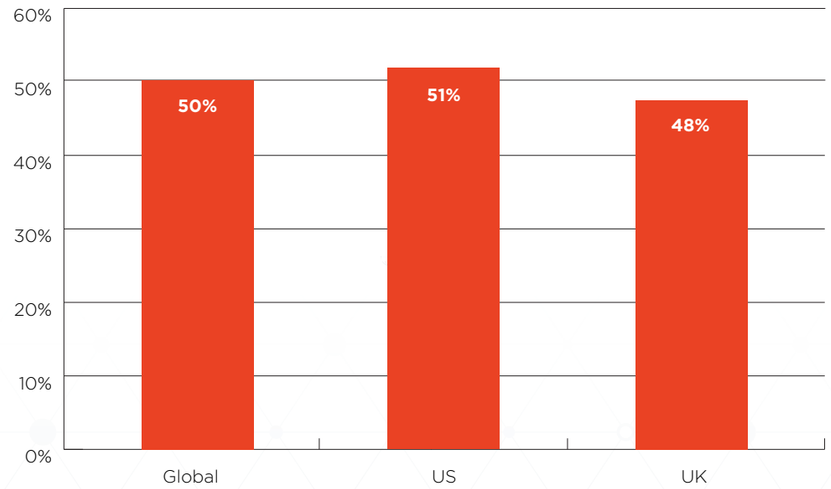

Totaling It Up:
The Cost of Manual Response vs. the Savings from Automation

Based on the data captured in the phishing incident response survey, we have all of the factors needed to estimate the cost of manually handling phishing incidents, average breach risk, and the potential cost savings of automating the process.

4.9 Hours Per Phishing Incident x 23,000 Incidents = 112,700 Hours of SOC Analyst Time

112,700 Hours ÷ 2080 FTE Hours Per Year = 54 FTEs

54 FTEs x $90,000 per FTE = $4.86M (£3.8M) Per Year

To calculate a custom ROI, visit **www.agari.com/roi**

A Massive Difference

Using averages for all variables, the detailed calculations above show a total annual cost to the SOC of $4.86 million and an average annual breach risk of $1.1 million—for a total cost $5.96 million per company. By implementing automated phishing incident response processes that reduce the time to triage, investigate, and remediate phishing incidents by 50%, organizations could save $4.37 million in SOC costs and $551,025 in breach risk—for a total savings of $4.92 million.

Customer Phishing and DMARC Trends

DMARC Confidential:
The Industry's Largest Snapshot of Adoption Rates Worldwide

Domain-based Message Authentication, Reporting and Conformance (DMARC) is an open standard email authentication protocol that helps businesses protect their brands and domains from being used to send fraudulent phishing emails. In a snapshot of 323 million Internet domains—the largest of any industry survey—we break down the state of DMARC implementation worldwide from October through December 2018.

Ditch the Domain Spoofing

DMARC gives brands control over who is allowed to send email on their behalf. It enables email receiver systems to recognize when an email isn't coming from a specific brand's approved domains and gives the brand the ability to tell the email receiver systems what to do with these unauthenticated email messages.

Failing to implement DMARC p=reject results in an easily identifiable vulnerability. Cybercriminals often spoof domains in order to send large volumes of spam, resulting in damage to the domain name's reputation, blacklisting, and even reputational damage to the brand name itself. The effects may first show up in complaints that outgoing emails aren't reaching recipients, often bouncing or being filtered by spam filters.

Brands looking to deploy DMARC are advised to start with DMARC p=none and work up to p=reject through a well-defined DMARC implementation plan. When enforcement policies are set properly, DMARC has been shown to drive down phishing rates impersonating brands to near zero.

For more information on DMARC and the benefits of adoption, visit **www.agari.com/dmarc-guide**

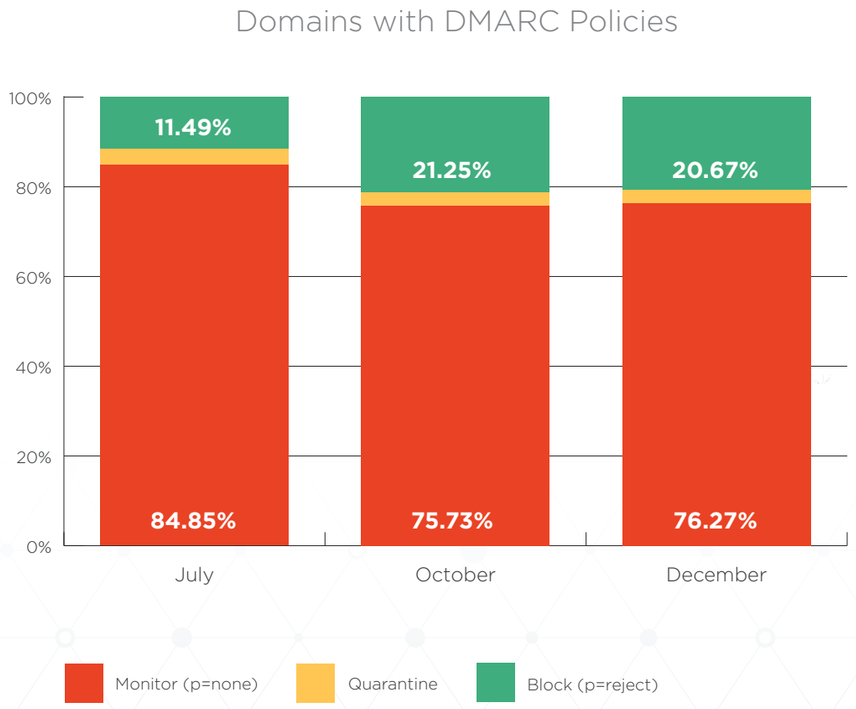

By crawling the entire public Internet domain space representing over 323 million domains—up from 283 million domains in our last report—ACID was able to generate a snapshot of DMARC implementation rates worldwide from October through December 2018. Overall, the DMARC adoption rate grew slightly in December. The pace of adoption slowed in December due to the holidays, but was up overall during the full fourth quarter of 2018.

KEY FINDINGS

- By January, ACID identified 6.1 million domains with valid DMARC records, up from 5.3 million in October. This represents modest growth of roughly 15% quarter over quarter.
- Factoring in automated actions of domain registrar-initiated DMARC records, the number of DMARC policies reduces to about 4.4 million.
- While the absolute volume of DMARC policies increased, so did the total universe of domains examined in our survey. Monitor-only continues to be the most common policy.

An Unprecedented View:
The Race to Increase Domains and Enforcement Levels

Each quarter, we set out to get a firm read on how vendors and DMARC service providers are helping organizations use DMARC to protect their domains from email impersonation scams. The size of our dataset offers an unprecedented view into the number of domains for which vendors have established DMARC records, as well as how many of those records have been set to the highest enforcement level of “p=reject.” This combination of data points offers a snapshot of market share and success rates for each of these vendors.

Vendor Scorecard

As a shorthand to determining a market share figure, we tabulated the number of times specific, well-known DMARC implementation vendors were specified as a recipient of reporting feedback via DMARC. The “rua” field that accepts an email address to receive aggregate DMARC data reports is a good proxy for this calculation. With this email address, the DMARC vendor typically accepts, parses, and visualizes the data on behalf of the customer. We included active vendors with more than 1,000 domains reported.

The following table shows a basic ranking of top vendors, corresponding to the number of domains that specify that vendor in the “rua” field. We then apply a second filter indicating the all-important percentage of domains at the highest possible DMARC enforcement policy setting (p=reject) for each vendor, which is the policy level that will block phishing messages.

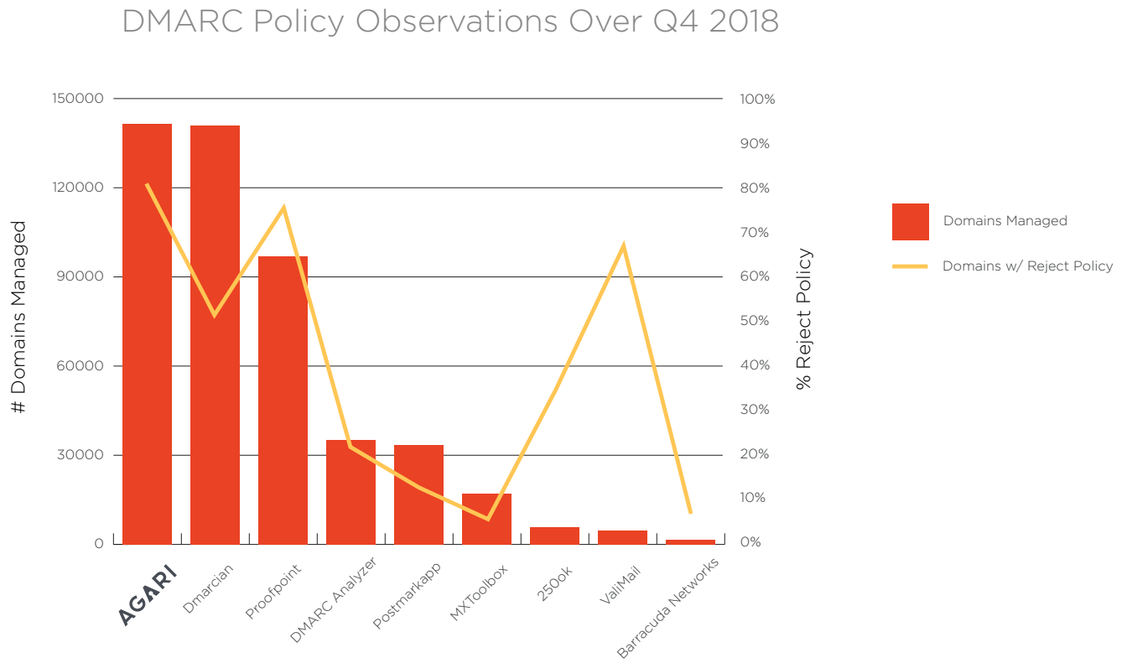

KEY FINDINGS

- The Sweet Spot: Category-leading vendors achieve that perfect combination of a large number of domains serviced across a wide range of industries matched with high levels of top enforcement policy implementation. Finding a company that has high marks in both is essential for those organizations looking to see success with DMARC implementation.
- Higher Quantities Can See Lower Enforcement: The “Goldilocks” ratio can be harder to achieve for mid-tier vendors, which tend to struggle with the radio of domains they service and what percentage of those records they succeed at converting to the highest enforcement policies. Category leaders with high numbers of enterprise clients can face this challenge as well, as it’s harder to have more enterprise domains set to reject.
- Quality Varies Wildly: About 500,000 of the domains that deployed DMARC are using a recognized DMARC provider, and about 2.8 million domains have DMARC deployed without using a major DMARC service provider. When selecting a vendor, enterprises with hundreds or thousands of domains should consider vendors that have both high numbers of domains and a high percentage of enforcement rate in order to better ensure success.

Q1 DMARC Global Sector Analysis:
Fortune 500

As we have done in the past, we looked at publicly available adoption data for the Fortune 500, Financial Times Stock Exchange 100 (FTSE 100), and Australian Securities Exchange 100 (ASX 100) to gauge adoption trends among prominent global organizations across geographies.

While the pace of DMARC adoption decelerated in the last quarter of 2018, the largest corporations around the world continue to gain traction in terms of email authentication. However, when considering the sizable proportion of "no record" and "monitor-only" policies, the current state of implementation at the start of 2019 is leaving customers, business partners, and brands vulnerable to phishing and the losses associated with email fraud.

Almost 85% percent of the Fortune 500 remain vulnerable to phishing, as are their customers. And while this is a 2% increase during the quarter, DMARC adoption remains dangerously low within the Fortune 500, enabling threat actors to exploit the considerable brand equity of even the largest, most well-known and most trusted companies in the United States.

DMARC Adoption – Nearly 50% of the Fortune 500 have yet to publish any DMARC policy. Nonetheless, this is a 2% improvement over just 90 days, and a marked improvement from 2017, when more than two-thirds of the Fortune 500 had no DMARC policy.

Quarantine Policy – Only 5% have implemented a quarantine policy to send phishing emails to the spam folder, about the same percentage as the previous quarter.

Reject Policy – One in 10 have implemented a reject policy to block phishing attempts impersonating their brands. This is up from just 8% from the previous report.

Fortune 500 DMARC Adoption

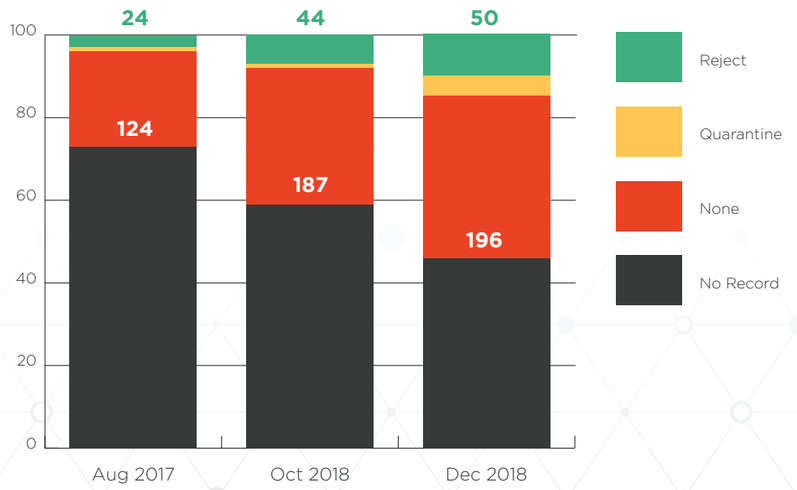

Q1 DMARC Global Sector Analysis:
FTSE 100

The Financial Times Stock Exchange 100 Index, more commonly known as the FTSE 100, is a share index of the top 100 companies listed on the London Stock Exchange (LSE) and is seen as the benchmark reference for those seeking an indication on the performance of the major companies listed in the United Kingdom.

Just as with their Yankee counterparts, the majority of the top 100 United Kingdom public companies do not have a DMARC record for their corporate domains. The lack of DMARC implementation dramatically increases the likelihood of the organization falling prey to not just fraud, but also a data breach, and all the reputational and financial damage that comes with it.

DMARC Adoption – Over the fourth quarter of 2018, there was a 3% increase in the number of FTSE 100 companies publishing a DMARC policy. While an improvement, that leaves 53% of these companies open to attack.

Quarantine Policy – Only one percent have implemented a quarantine policy to send phishing attempts to spam. This percentage is unchanged from last year.

Reject Policy – Only 11 companies have implemented a reject policy to block phishing-based brand impersonations. That's a 2% increase from the previous period.

FTSE 100 DMARC Adoption

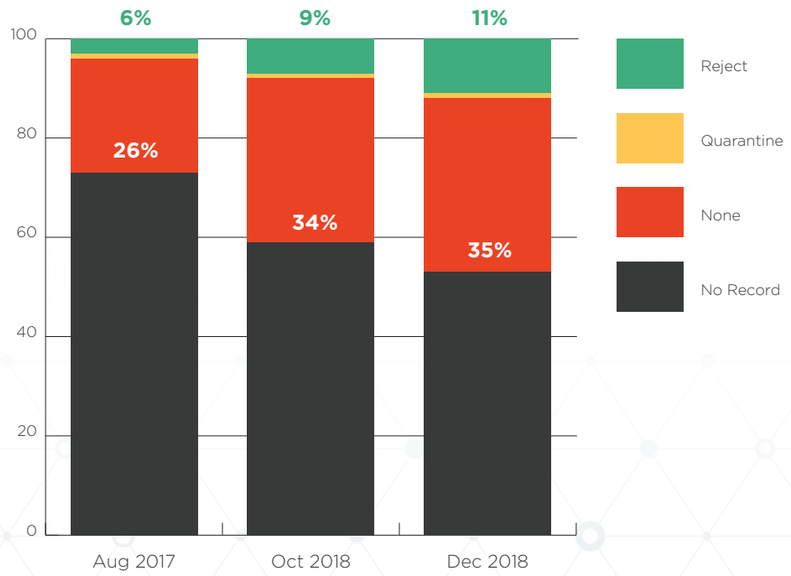

Q1 DMARC Global Sector Analysis:
ASX 100

The ASX 100 is Australia's stock market index, representing its top 100 large and mid-cap securities.

Fewer than half of ASX companies have taken, at a minimum, the first step in adopting DMARC to combat the threat of phishing attacks bearing their name. Clearly, considerable educational initiatives are needed to increase DMARC adoption in this region.

DMARC Adoption – More than half of the ASX have yet to publish any DMARC policy.

Quarantine Policy – Two percent have implemented a quarantine policy, marking an uptick from 1% in the previous quarter. That said, this is only an increase of one organization, showcasing how few companies are thinking about email security.

Reject Policy – Only seven percent have implemented a reject policy—the same as the prior quarter.

ASX 100 DMARC Adoption

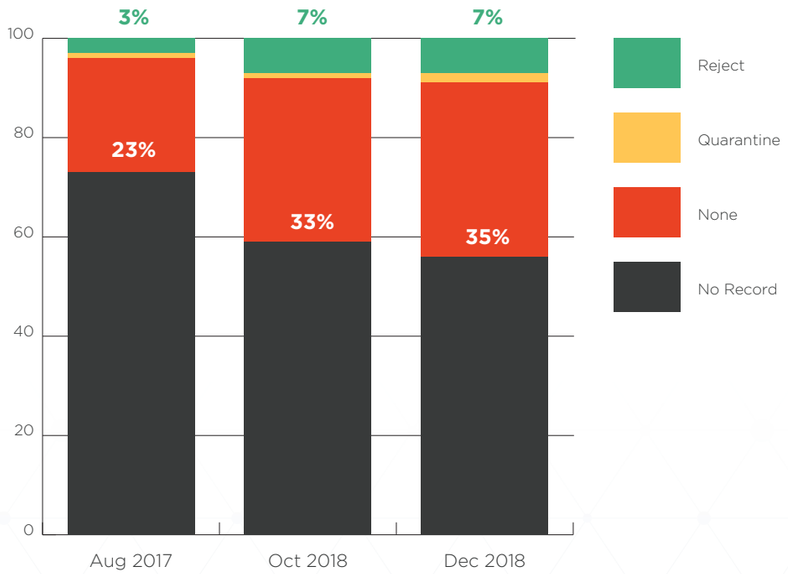

Q1 Large Sector Analysis:
US Government Maintains Its Lead

As part of our analysis of DMARC adoption, we examine public DNS records for primary corporate and government website domains of large organizations with revenues above $1 billion.

As the chart below shows, when viewed from a DMARC policy attainment perspective, the US Government is hands down the DMARC leader across all major sectors. Driven by an executive branch security mandate implemented over the past year, a stunning 81% of domains have implemented DMARC at a p=reject, or block, enforcement policy—up from 76% in a single quarter.

DMARC Policy & Enforcement Trends for Key Industries

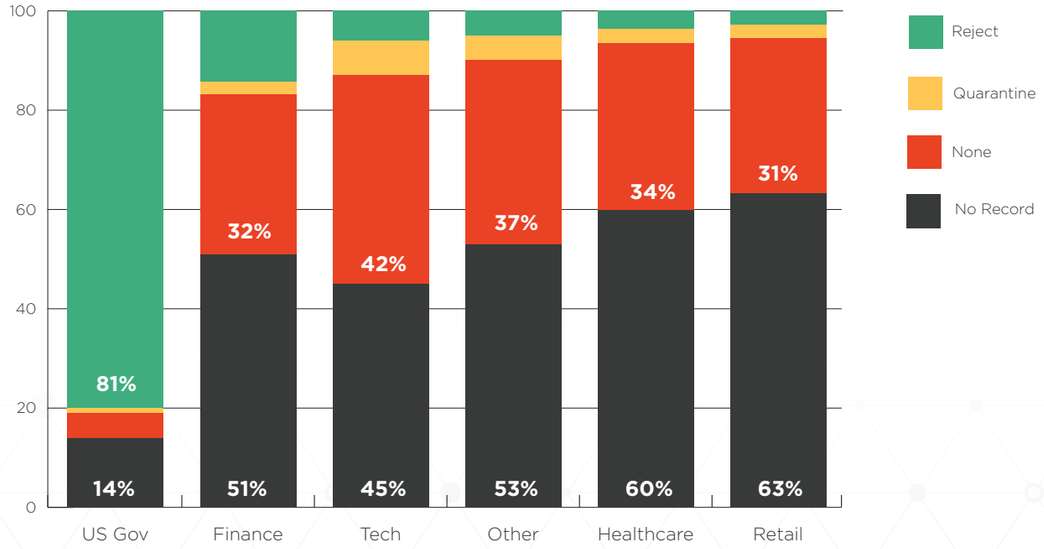

Q1 Industry Enforcement Comparison:
The Agari Advantage by Vertical

A look at how enforcement rates across industries compare with those of Agari customers, according to data from the Agari Email Threat Center.

Aggregating real-time DMARC statistics from the domains of top banks, social networks, healthcare providers, major government agencies, and thousands of other organizations, the Agari Email Threat Center is the largest set of detailed DMARC data in the world based both on email volume and domains. To generate real-time threat intelligence, the Agari Email Threat Center analyzed more than 583 billion emails over 18,729 domains from October through December 2018.

Segmenting by the same industry groupings presented in the previous section, we compare the respective enforcement levels for each vertical category with that of Agari customers. Consistent with overall industry dynamics, the government sector (heavily biased toward US government) continues to dominate Threat Center rankings. Following government, healthcare has edged out the technology sector as the next-highest ranked vertical for the percentage of domains at enforcement.

This is notable, as healthcare as a vertical moved from the lowest enforcement rate in the Threat Center in Q4 2017 to rank second by year-end 2018. This momentum is likely driven by the National Health ISAC, which issued a companion pledge for DMARC attainment to match that of the US Government's Binding Operational Directive (BOD) 18-01. BOD 18-01 was issued in October 2017 and has been the driving factor behind the sky-high adoption rates for executive branch agencies.

Percentage of Domains at Enforcement

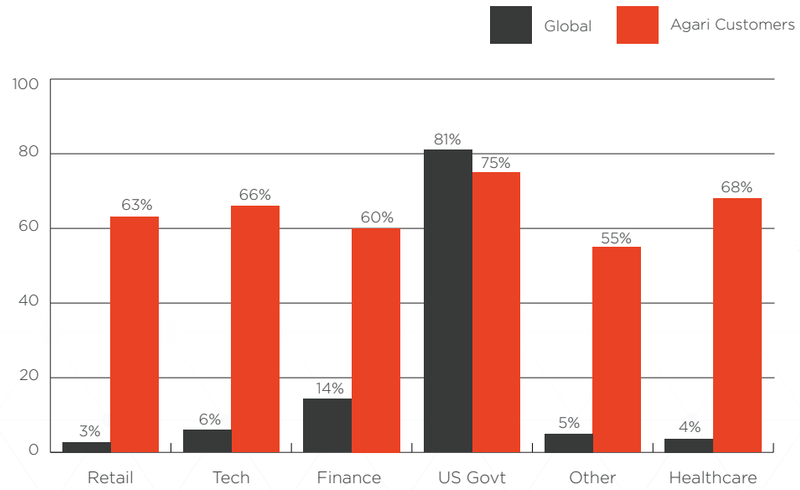

Note: The Threat Center tracks authentication statistics across active domains belonging to Agari's customers. Passive or defensive domains that don't process email will not be reflected in the totals. Overall, as indicated previously, the Agari reject rate across all industries in the global domain snapshot is 82%.

BIMI Builds Momentum:
Taking DMARC One Step Further

Brand Indicators for Message Identification (BIMI) is a standardized way for brands to publish their brand logo online with built-in protections that safeguard the brand, application providers, and consumers from impersonation attempts. BIMI-enabled logos be easily incorporated into messaging and social media applications.

For instance, a retail brand can use BIMI to display its logo next to its messages, enhancing its brand presence as well as providing assurance to recipients that the message is safe to open. BIMI will work only with email that has been authenticated through DMARC standard and for which the domain owner has specified a DMARC policy of enforcement, so only authenticated messages can be delivered.

Q1 BIMI Snapshot: A 69% Increase in Brand Adoption

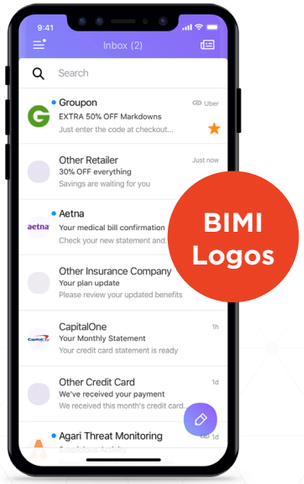

About This Report

This report contains metrics from data collected and analyzed by the following sources:

Employee Phishing and BEC Data

For inbound threat protection, Agari uses machine learning—combined with knowledge of an organization's email environment— to model good or authentic traffic. Each message received by Agari is scored and plotted in terms of email senders' and recipients' identity characteristics, expected behavior, and personal, organizational, and industry-level relationships. For the attack categorization analysis, we leveraged anonymous aggregate scoring data that automatically breaks out identity deception-based attacks that bypass upstream SEGs into distinct threat categories, such as Display Name Deception, Compromised Account, and more.

Phishing Incident Response Trends

This report presents results from a custom survey conducted by Agari during Q4 2018. The following charts summarize the demographics and location of the respondents.

Respondent Characteristics

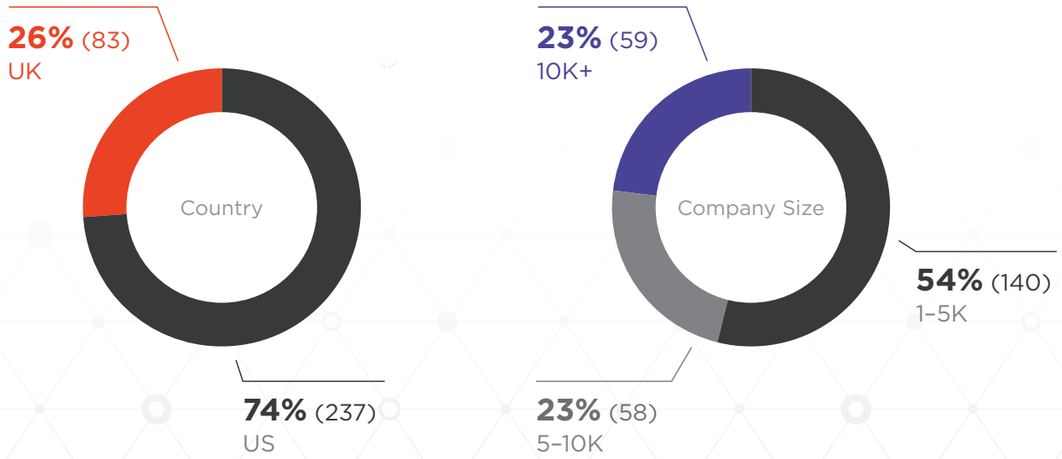

Global DMARC Domain Analysis

For broader insight into DMARC policies beyond what we observed in email traffic targeting Agari's customer base, we obtained and analyzed hundreds of millions of domains over the course of Q4 2018. This overall set represents virtually all the publicly accessible domains in DNS over the course of Q4. At the end of the quarter, we crawled 323,245,038 domains, ultimately observing 6,126,323 with recognizable DMARC policies attached.

Quarter over quarter, our base domain list increased by over 40 million, mostly in newly detected country code top-level domains (CCTLD). This constantly growing list of domains serves as the basis for trend tracking in subsequent reports.

About the Agari Cyber Intelligence Division

The Agari Cyber Intelligence Division (ACID) is the only counterintelligence research team dedicated to worldwide BEC and spear-phishing investigation. ACID supports Agari's unique mission of protecting communications so that humanity prevails over evil. ACID uncovers identity deception tactics, criminal group dynamics, and relevant trends in advanced email attacks. Created by Agari in 2018, ACID helps to impact the cyber threat ecosystem and mitigate cybercrime activity by working with law enforcement and other trusted partners.

Learn more at acid.agari.com.

About Agari

Agari is transforming the legacy Secure Email Gateway with its next-generation Secure Email Cloud™ powered by predictive AI. Leveraging data science and real-time intelligence from trillions of emails, the Agari Identity Graph™ detects, defends, and deters costly advanced email attacks including business email compromise, spear phishing and account takeover. Winner of the 2018 Best Email Security Solution by SC Magazine, Agari restores trust to the inbox for government agencies, businesses, and consumers worldwide.

Learn more at www.agari.com.

AGARI CYBER
INTELLIGENCE DIVISION

Discover How Agari Can Improve Your Current Email Security Infrastructure

As your last line of defense against advanced email attacks, Agari stops attacks that bypass other technologies—protecting employees and customers, while also enabling incident response teams to quickly analyze and respond to targeted attacks.

Get Free Trial **www.agari.com/trial**

Visit the Agari Threat Center

To see up-to-date global and sector-based DMARC trends across the Agari customer base, visit: **www.agari.com/threatcenter**

Calculate the ROI of Implementing Agari

To discover how much money you can save by adding Agari to your email security environment, visit: **www.agari.com/roi**

AGARI