agari
by HelpSystems

How to Conquer Targeted Email Threats: SANS Review of Agari Advanced Threat Protection™

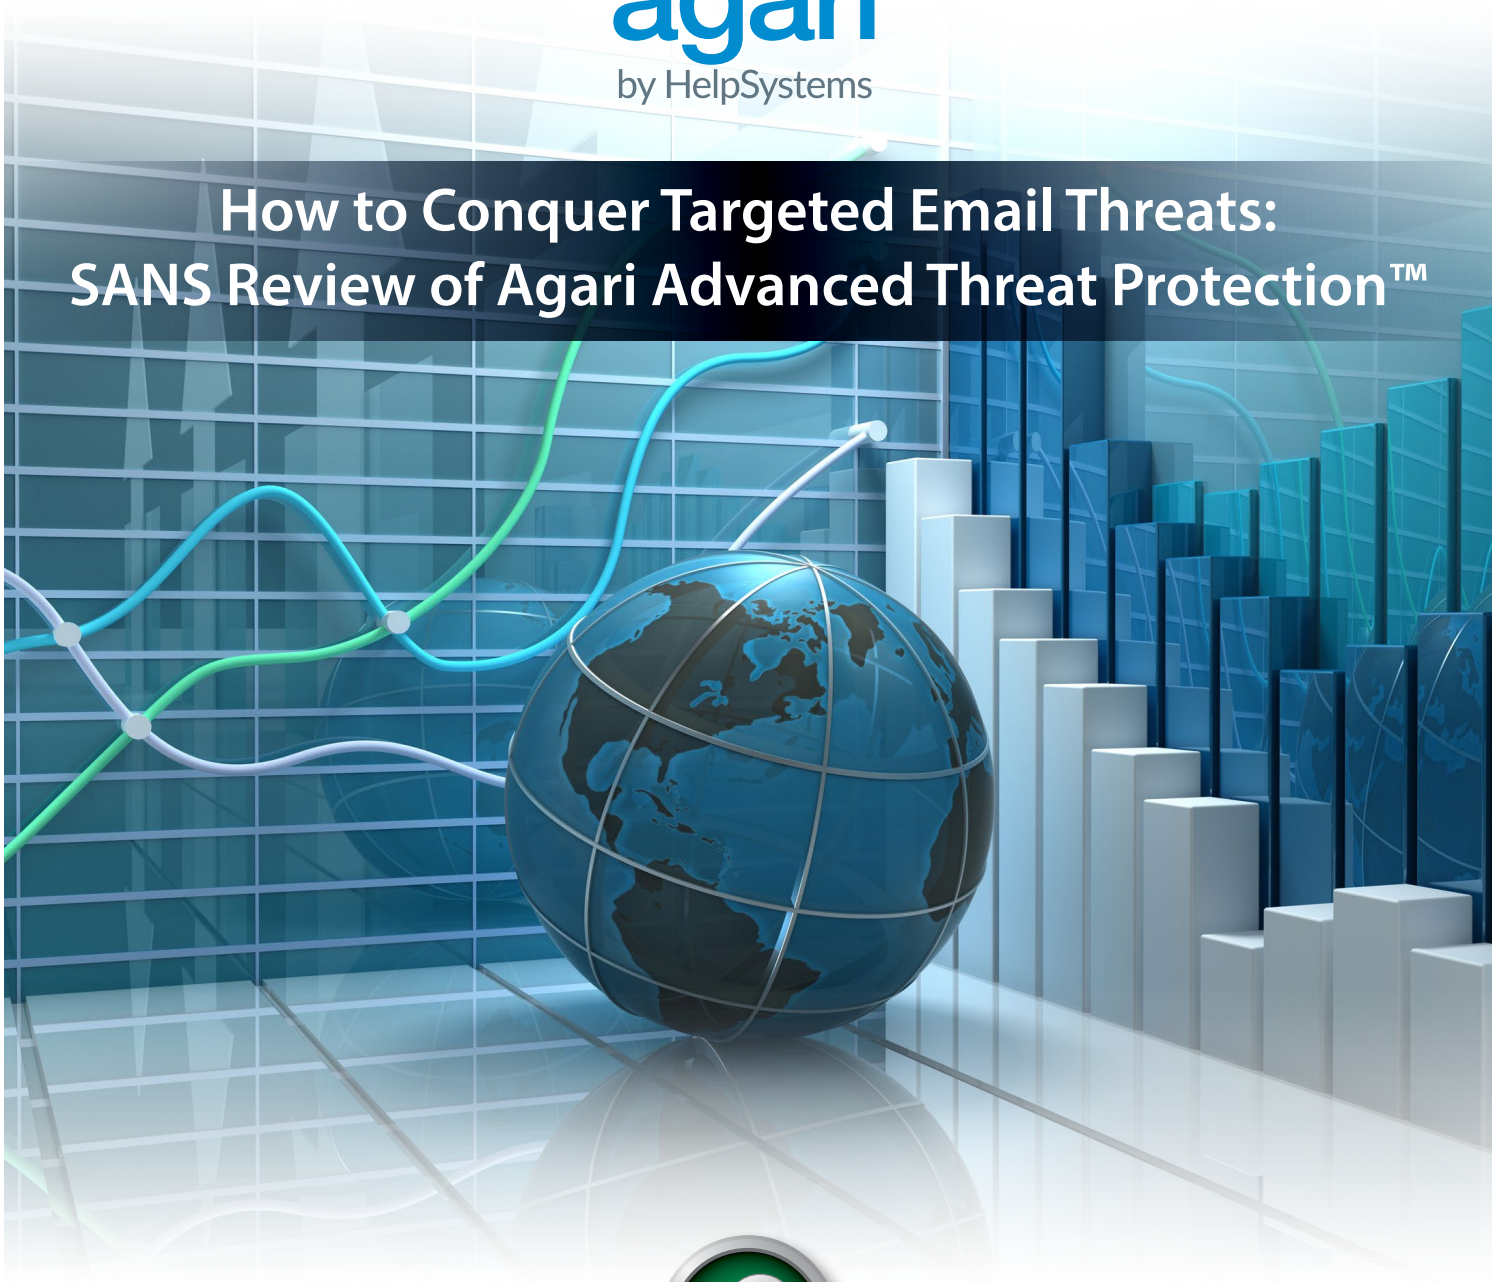

A SANS Product Review

Written by Dave Shackleford

May 2017

Sponsored by
Agari

Introduction: Email Is a Major Threat Vector Today

As attackers have continually sought ways around organizations' defenses that will avoid detection, they have made social engineering of end users the basis for most of today's targeted threats.[1]

For instance, in the 2016 SANS survey on endpoint security, three quarters of respondents who had identified impactful threats to their enterprises said the threats entered the environment via email attachments. Nearly half (46 percent) of impactful attacks were executed when users clicked links in emails. In another SANS survey, which was focused on the financial sector, 50 percent of respondents cited spearphishing and "whaling" (phishing focused on high-value targets) as their top threat vector.[2]

From Email to Infection in Under Five Minutes

Verizon's 2016 Data Breach Investigations Report (DBIR) reveals that of the 9,576 incidents investigated, where email was used to send malware, 30 percent of the messages were opened and links or attachments clicked within five minutes of targets receiving the malicious email.[3]

Why are our traditional email and endpoint security tools failing us? First, most email deployments lack any authentication of outside senders. Given this vulnerability, it's trivial to execute spoofing and falsified email content that purports to come from a trusted entity the recipient knows and trusts. Second, attackers are using cloud-based email and "detection-busting" techniques such as fake identities, deceptive sender names and phony domains to beat defenses.

Clearly, given the prevalence of email-borne threats, protecting email infrastructure and end users needs to be a high priority for all security teams today. To this end, SANS had the opportunity to review Agari Advanced Threat Protection™ and the Agari Secure Email Cloud™.

Our review found that Agari's sophisticated trust analytics capably identified suspicious and malicious email content and also provided an easy-to-use policy engine that can be used to flag, block and quarantine email content as needed. The interface was simple to navigate, and enterprise defenders and operations teams will have no trouble quickly controlling some of the major threats facing them through email today.

This paper looks at Agari's approach to email threats, providing a high-level overview of Agari Advanced Threat Protection and our findings on how well it safeguards email.

Note: Agari Advanced Threat Protection™ and Agari Secure Email Cloud™ are trademarks of Agari.

[1] "Exploits at the Endpoint: SANS 2016 Threat Landscape Survey," September 2016, www.sans.org/reading-room/whitepapers/firewalls/exploits-endpoint-2016-threat-landscape-survey-37157

[2] "From the Trenches: SANS 2016 Survey on Security and Risk in the Financial Sector," October 2016, www.sans.org/reading-room/whitepapers/analyst/trenches-2016-survey-security-risk-financial-sector-37337

[3] Verizon's 2016 Data Breach Investigations Report, www.verizonenterprise.com/verizon-insights-lab/dbir/2016

Agari's Approach to Email Threats

Targeted email attacks account for 95 percent of breaches today, according to Verizon's latest data breach report. Attackers are leveraging sender identity deception as a core attack technique, one that is able to evade standard secure email gateways and advanced threat protection solutions that address identity deception only indirectly (for example, by searching for malicious content or identifying an email containing a known bad URL).

With Advanced Threat Protection, Agari brings something different to the table. Agari focuses on identity deception as the starting point, especially in phishing, business email compromise (BEC) and related scenarios. The premise is that advanced email threats use some form of identity deception to attack the weakest part of our defenses, which is human perception and judgment. By automating the process of analyzing the trust and authenticity of all email, Agari Advanced Threat Protection can more effectively protect against targeted email attacks that will defeat existing defenses that rely on detecting malicious content or behavior. To this end, Agari generates an actionable "Trust Authentication Model" that encompasses several core concepts:

- **Global visibility.** Agari has exposure to a vast array of email and domain information globally, and it continually builds and evaluates a "trust score" based on the data it sees.
- **Machine learning and data science.** Agari's cloud-based trust network builds continuous models of behavioral patterns that leverage real data science and use machine learning to help build more accurate models of tailored threats to customer email domains and accounts.
- **Trust and authenticity.** Unlike most email protection tools, Agari uses a combination of analytics to score and apply policies to unauthenticated email coming from any domain (trusted partners, customers, vendors) and Domain-based Message Authentication, Reporting and Conformance (DMARC) protection for domains that publish a DMARC policy. With Agari, organizations can publish a DMARC policy for their own domain, which prevents all direct spoofing of that domain and stops the CEO-to-CFO wire transfer scams and human resources W-2 scams immediately with 100 percent effectiveness.

DEFINITION

Domain-based Message Authentication, Reporting and Conformance (DMARC): Builds on both the Sender Policy Framework (SPF) and Domain Keys Identified Mail (DKIM) standards to help prevent spoofing of the "header from" address in email. DMARC does this by first matching the "header from" domain name with the "envelope from" domain name used during an SPF check, then matches the "header from" domain name with the "d= domain name" in the DKIM signature.

The Many Faces of Email Threats

A key element of deception used in email threats is spoofing, the technique of forging email headers to dupe recipients into believing they are receiving messages from someone other than the actual source. Criminals use email spoofing in phishing and spam campaigns because people are more likely to open an email when they think it has been sent by a legitimate source. Along with spoofing, attackers use many other deceptions to hide malicious emails from traditional defenses, including the following:

- **Social engineering—**Deceptive messages—typically email messages—that deliberately entice an end user to take some action that would cause harm, such as clicking on an attachment or link, sending sensitive information or transferring funds
- **Phishing—**A form of identity deception typically involving impersonation of a trusted brand where the goal of the criminal is to trick the recipient into entering credentials on a site the criminal controls
- **Spearphishing—**Similar to phishing but with a much more specific and targeted set of content that is tailored specifically to the email recipient
- **Whaling—**Spearphishing that targets senior executives or other extremely high-value recipients

> One of the best-known examples of whaling is the email sent to John Podesta during the 2016 presidential election that appeared to come from Google, asking and asked him to enter his password on a malicious site.

- **Ransomware—**Malware (often distributed via email) that hijacks the end user's computer and demands a ransom be paid
- **Business email compromise—**An email attack that uses identity deception, typically involving a trusted colleague or partner, where the goal of the criminal is to trick the recipient to send either funds or sensitive data

Review Environment

For our review, SANS worked in Agari's live, internal production environment. In this environment, Agari blocks thousands of malicious emails daily. Although we did not install an Agari sensor on premises, we did review the installation and operations documentation, and the process seems simple and straightforward.

Agari readily supports both in-house (for example, Microsoft Exchange) and hosted (for example, Office 365 or G-Suite) email deployments and uses a cloud-based monitoring and review infrastructure for its customers. Once the sensor is configured (either as a separate installed platform in the data center or as a virtual sensor in the Agari cloud) and email is being sent to the sensor, Agari can start applying policy rules and protective controls to the email for any client.

Sensor Configuration

Within the Agari console, we first reviewed the sensor configuration for our test environment. This was found by clicking Manage/Sensors in the console, which showed us the "in-cloud" sensor used by the Agari production environment. See Figure 1.

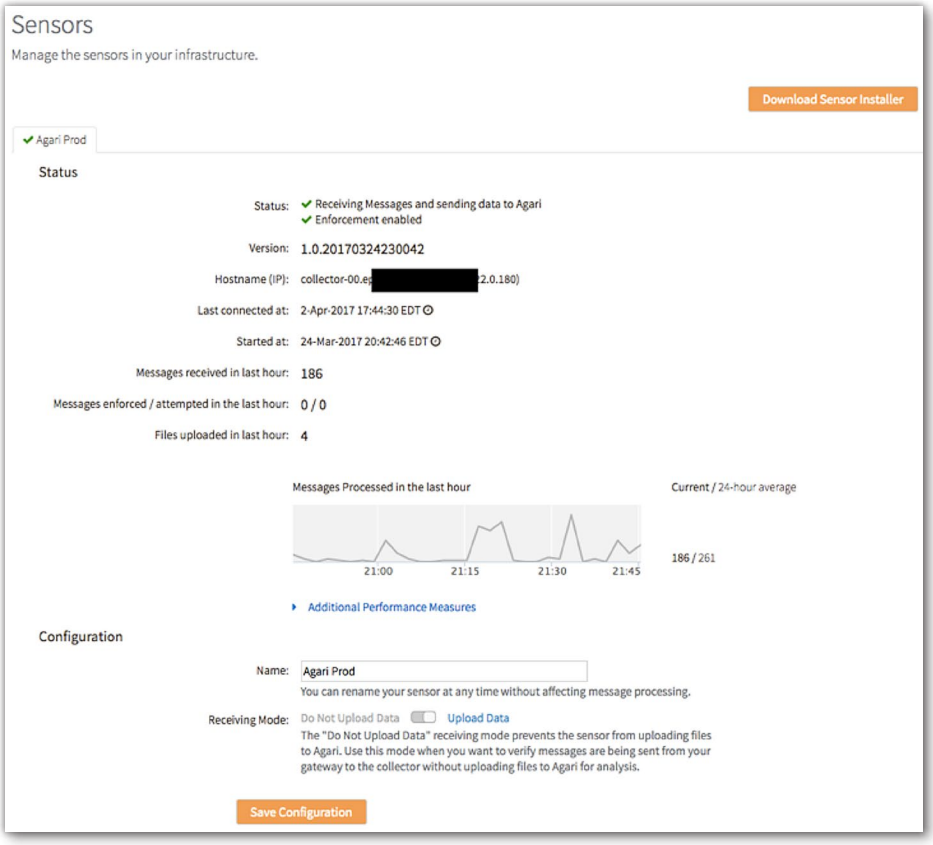

Figure 1. The Agari Email Sensor

Aside from name and basic file-sharing features, we had to configure very little within the sensor.

Creating User Accounts

We then looked at the different types of users we could create for administration and access. By clicking Manage/Users, we could readily see the list of existing users for the organization, as well as create new users that fall into several defined categories:

- **Organization administrator—**An administrator who can modify the organization's information and settings only
- **User administrator—**An administrator who can access and modify user information
- **Auditing user—**A user who can view user data and all audit trails of activity
- **Read-only user—**A user who can simply read information about the email activity, as well as defined reports

The user administrative page (found after clicking "Add user") is shown in Figure 2.

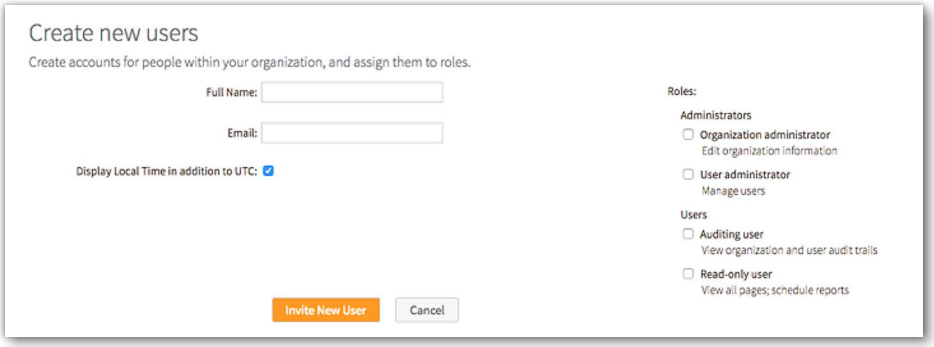

Figure 2. Adding a User to Agari

It was also simple to create specific groups of users that we would then reference in policies. By navigating to the Manage/Address Groups section of the console, we reviewed the existing address groups in place for Agari's environment, which includes Engineering, Email Enforcement, Board Members and Advisors, and Executives, as shown in Figure 3.

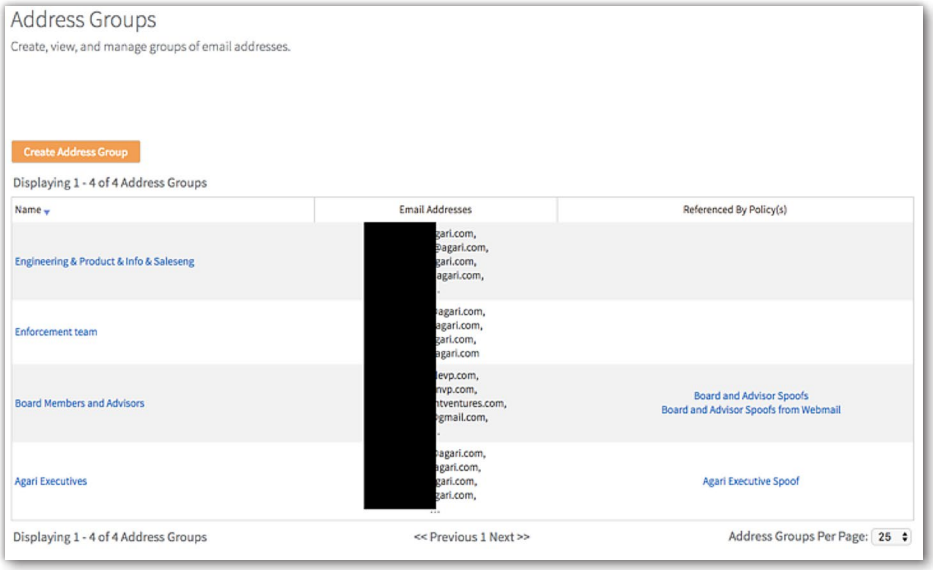

Figure 3. Agari Address Groups

We also created a simple address group for several personal and business emails that could be identified as trusted. See Figure 4.

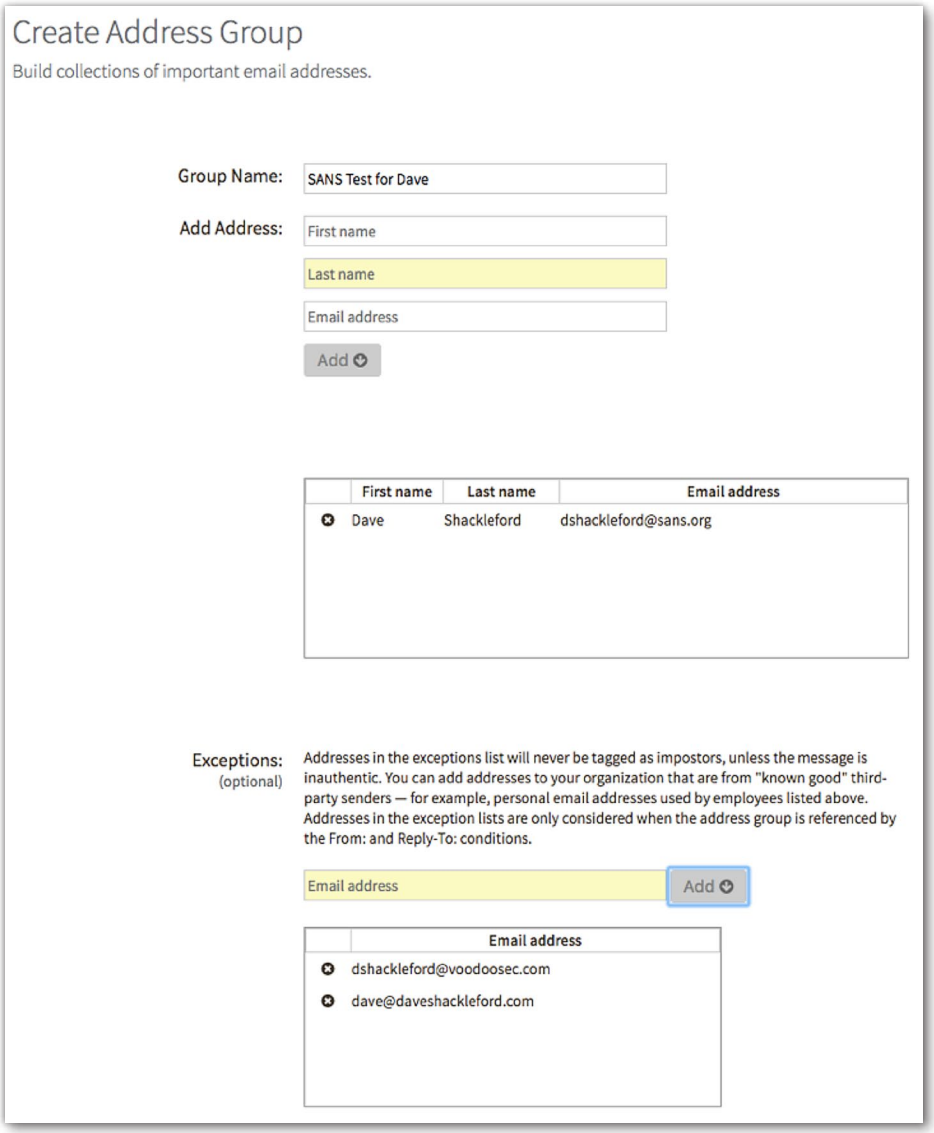

Figure 4. Creating a Simple Address Group

TAKEAWAY:

Creating users, groups and role-based access models for Agari Advanced Threat Protection was straightforward and provided the necessary level of granularity to ensure all types of security and operations teams have the access they need.

Adding personal emails that are also affiliated with trusted users is important because it reduces the number of unwanted blocking of messages to and from these addresses, and it catches phishing attempts that other filters have missed; for example, emails that use a different email domain than the one associated with the main organization.

Finally, we reviewed some of the settings for customer organizations, which included basic contact settings, enforcement capabilities with a defined email subject line (in case of quarantine), and sensor and user settings that an admin could configure (login and password settings, for example). All of these were simple to set and manage, although we did not configure them specifically during this test.

Flexible, Far-Reaching Policies

With the basic organizational elements in place (users and groups), we now turned to exploring Agari policies. Before defining policies, though, we looked at domains, tags and IP addresses. Agari allows you to define domains in categories with tags. Under the Analyze/Domains section of the console, we saw a list of any domains that have sent mail to our organization. These domains can be classified with a number of tags set up by the Agari team. These tags include the following:

- Internal
- Service
- Partner
- Customer
- Webmail
- Social
- Consumer
- Marketing

The Agari team can set up other customized tags, as well, but those listed above are some of the most common domain tags currently used by most administrators.

Valuable Domain Information

From the console, we clicked on a domain to see information about it, including its reputation score (determined by Agari using DMARC, Sender Policy Framework and other analytics), message volume and IP addresses associated with the domain, as shown in Figure 5.

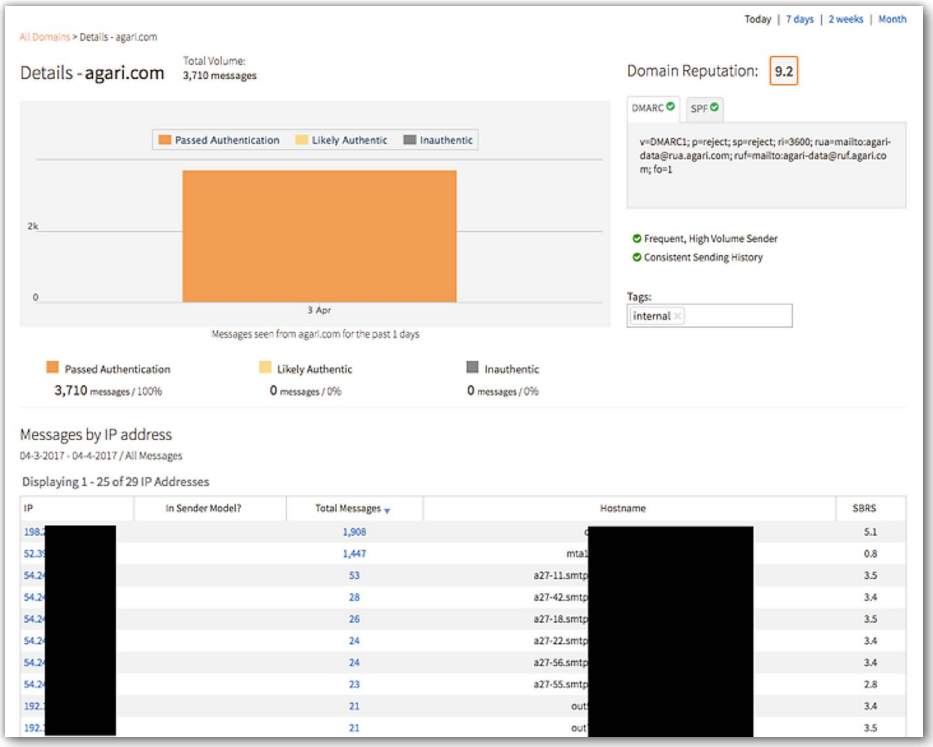

Figure 5. Domain Information in the Agari Console

Similarly, we also reviewed IP addresses that have been seen by Agari by clicking Analyze/IP Addresses.

IP Address Drilldown

Although it isn't possible to assign tags to the IP addresses, the ability to dive into details associated with these addresses is similar to the domains review. By clicking on an IP address, we could review any associated domains, message counts from the address and general authenticity scores, and we could even perform WHOIS lookups directly in the console. Details for an Agari IP address are shown in Figure 6.

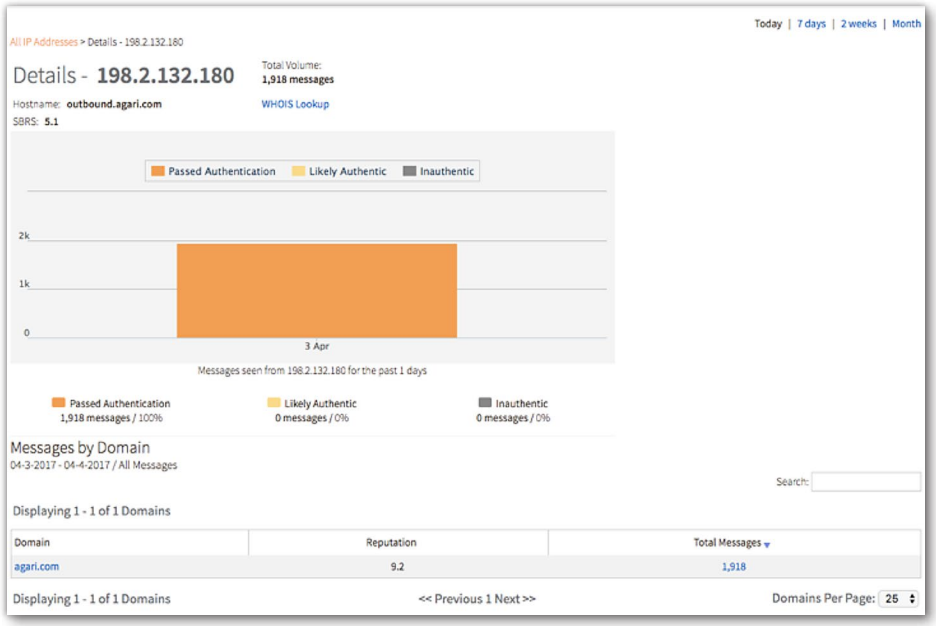

Figure 6. IP Address Details in the Agari Console

With policies configured, Agari compared when certain fields don't match (when they should), indicating likely spoofing attempts.

Policy Management

After we explored the domains and addresses seen coming to the Agari sensor, we explored email policy settings and monitoring. Under Manage/Policies in the console, we found four main sections.

- The System Notifications section allows a security or operations admin to be notified if a sensor stops functioning, if a sensor is removed or if other operational issues occur.
- The primary area to configure is the Policies section. The Agari team has several policies already in place that we explored, and we created a new policy from scratch.
- After giving the policy a name, the user can configure several sections of the policy in the Configure Content Matching section, including content matches in traditional email header fields, such as From, Reply-To, To and Subject.
- In the Specifying the Sending Domain section, domain tags can be included.

Figure 7 shows the policy creation interface.

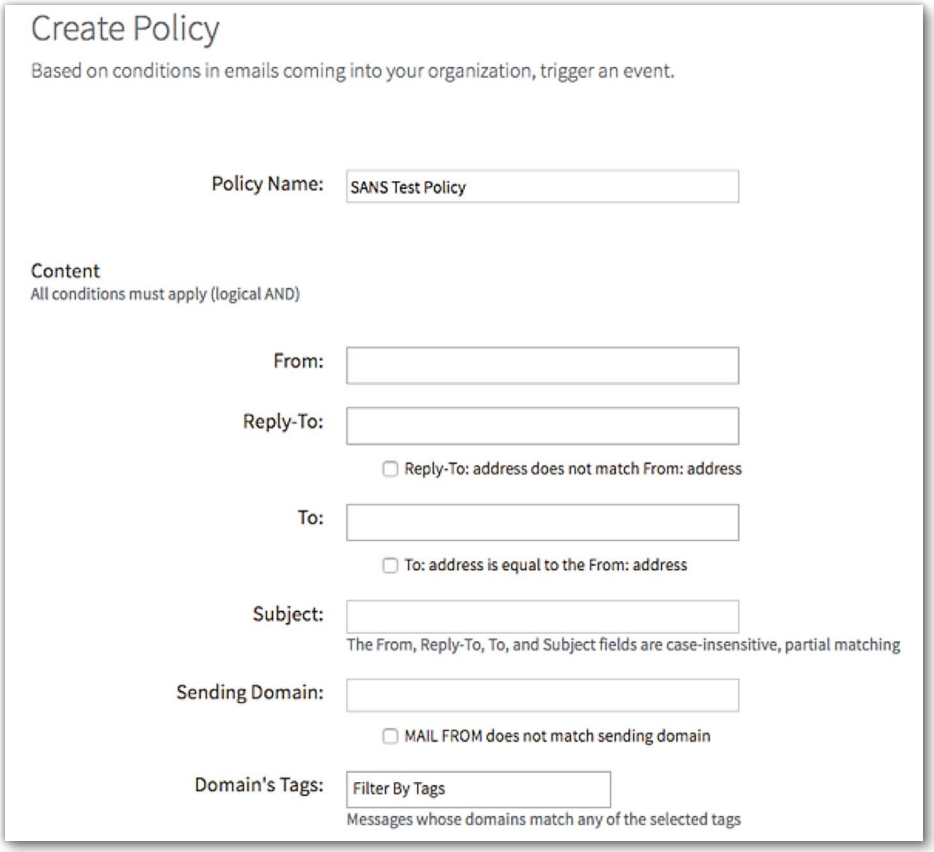

Figure 7. Basic Policy Detection Options

Trust Scoring Rules

The next set of policy options is more specific and focuses on Agari's scoring for trust (overall), authenticity, domain reputation and sender-based reputation scores (SBRSs). Filtering specific IP addresses is also possible. These policy rules allowed us to create a much more specific set of "grouped" rules that fell into ranges more or less trustworthy overall. See Figure 8.

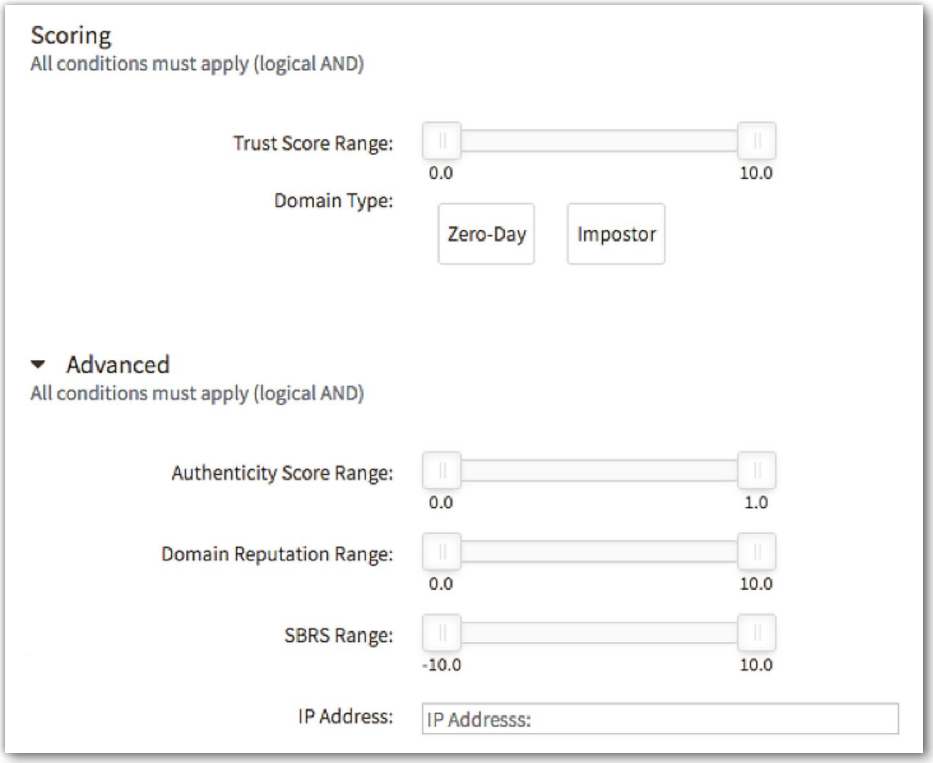

Figure 8. Agari Policy Scoring Rules

Defining Actions to Take

Finally, policy actions are defined. First, policies can either be "enforced" or not, which simply means they are moved to a quarantine, trash or deleted-items folder or allowed to pass to the recipient.

The next options relate to notification and who gets notified and when. A "digest" of notifications can also be sent when a certain threshold of alerts is generated within an hour. See Figure 9.

TAKEAWAY:

The Agari team has a variety of well-designed policies already in place, including "Untrusted + enforced" (all messages are blocked if they have a trust score between 0.0 and 1.1), "Untrusted" (with a trust score between 1.1 and 2.1), and several policies focused on executive and board member spoofing (where the From field matches the internal addresses but comes from outside).

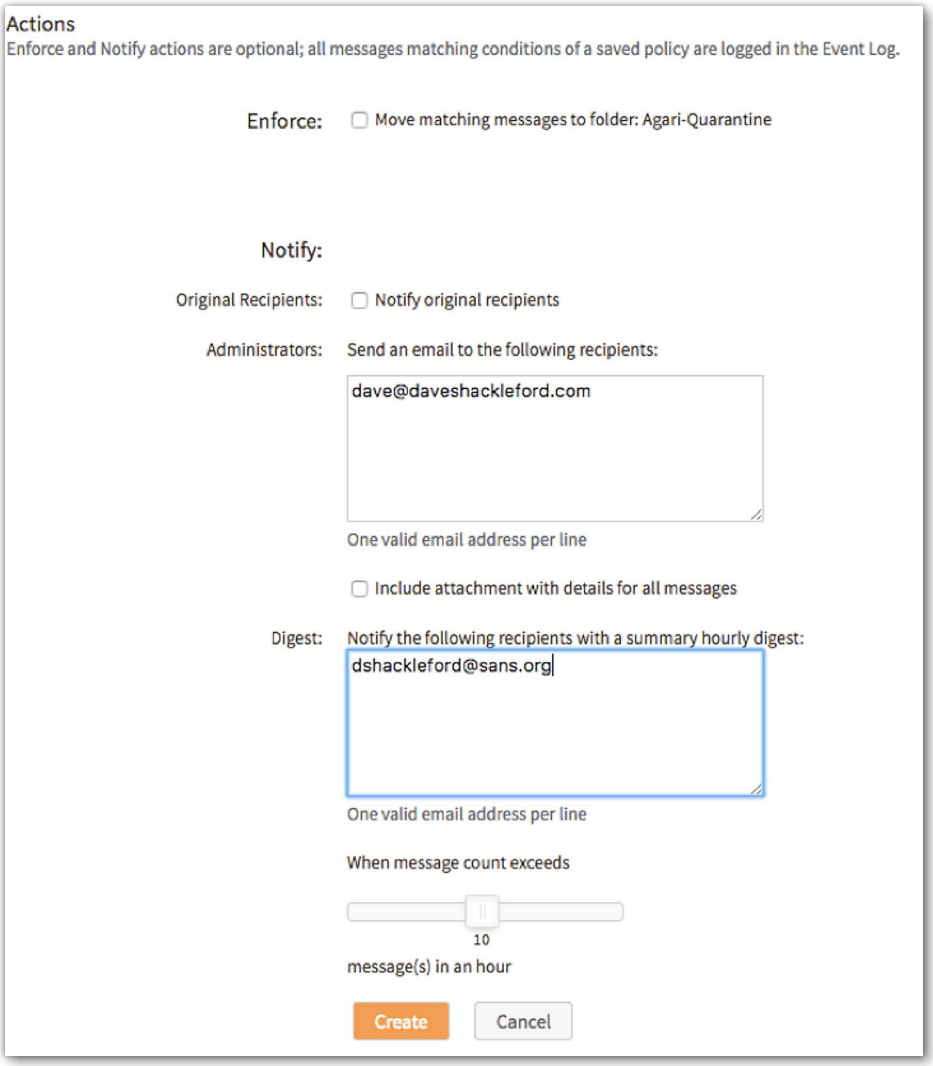

Figure 9. Policy Notification Options

That's all there is to generating policies within the Agari interface.

Creating Audit Trails

The final category of policy administration is the Event Log, where policy events are chronicled and listed for security audit and review. Policy events and system notifications can be reviewed over a period of one day, seven days, two weeks or a month. See Figure 10.

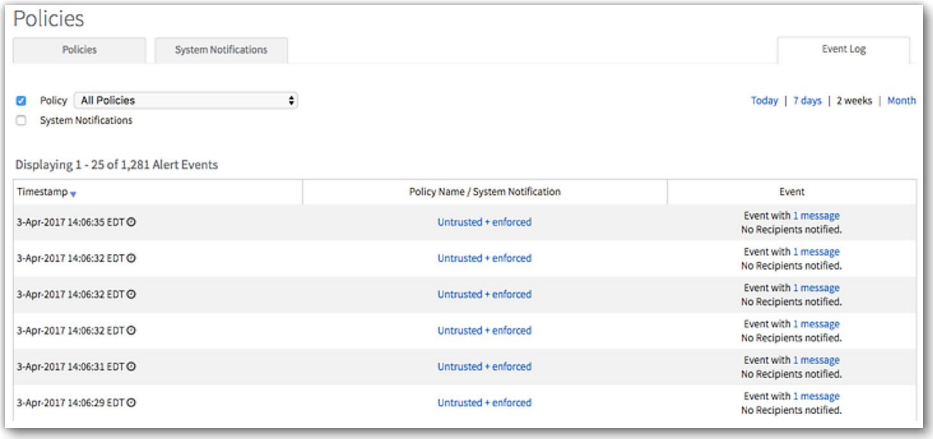

Figure 10. Agari Policy Event Log

To review the effectiveness of Agari's policies (numerous policies are in place within a live production environment), we started with the Risk Overview dashboard. The graph in Figure 11 maps the messages and events logged within the organization's Agari deployment to Authenticity Score and Domain Reputation. This gives defenders a quick and simple view of the noted messages and events coming into the organization, starting with a simple volumetric view of the top sending domains.

Figure 11. Agari Risk Overview: Top Sending Domains

This dashboard proved to be enormously flexible, showing specific types of attacks and whether these perceived attackers were merely untrusted (based on Agari's risk analytics scoring) or also suspicious (based on threat intelligence or specific behavioral attributes, attachments, etc.). At a glance, we were able to focus on a broad variety of trusted and untrusted email content seen by Agari.

Scoping a Zero-Day

For example, we drilled into zero-day attacks (from known malicious domains) and highlighted one of the domains listed (salesdatainfo.com). There were 37 events related to this domain in the previous two weeks (out of a total of 267 messages of type "zero day"), as shown in Figure 12.

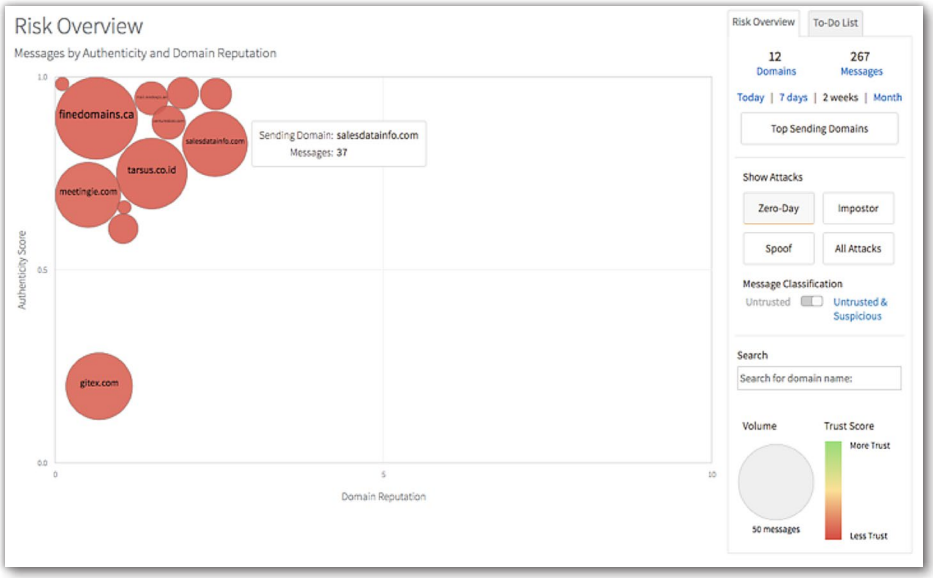

Figure 12. Zero-Day Email Events Caught by Agari

This dashboard can be easily used to drill into more detail simply by clicking any of the domains/events noted within.

Getting the Message

This now brings us to the Message Search feature, which we utilized to look up very specific message attributes, including date ranges, trust scores, domains and IP addresses. See Figure 13.

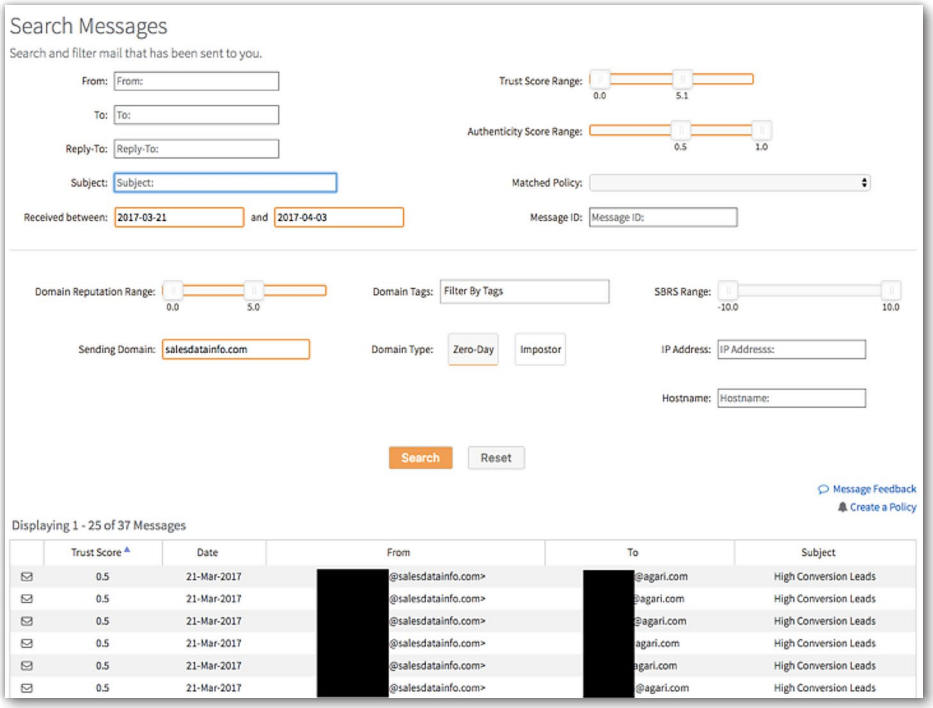

Figure 13. Granular Agari Query for a Specific Zero-Day Domain

Figure 13 demonstrates a simple search for our "salesdatainfo.com" domain as the sender for a two-week period of time. This search returned all the emails found in that time period. By clicking on any of the email results (each row is clickable), we can see the detailed view of the email header content, as shown in Figure 14. This content explains that Agari flagged the content as likely coming from a zero-day domain (has not been seen before as a sender or through email threat intelligence), leading to a very low trust score.

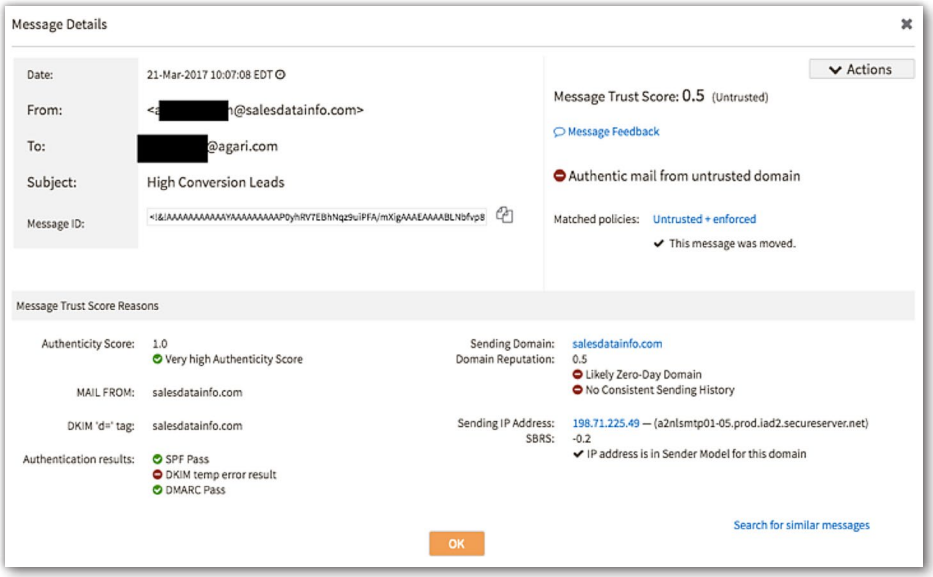

Figure 14. Email Detail from Attack/Event Queries

The full email can also be downloaded, forwarded to someone/somewhere else, viewed or turned directly into a policy from the "Actions" menu.

Detecting a Whaling Attempt

We also utilized the Analyze/Messages menu option to get a broad view of the Untrusted, Suspicious and Trusted messages seen over a specific time period, along with the list of messages in the chosen category. See Figure 15.

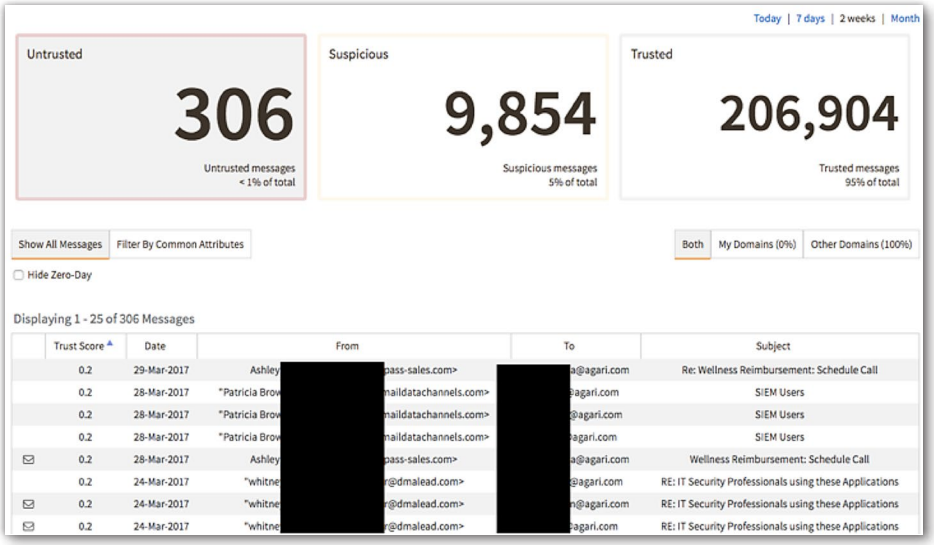

Figure 15. Agari Messages Query Dashboard

We drilled into an untrusted message that had a very low trust score and looked to be a spearphishing message. This was sent to an Agari board member and had a bogus domain source and a subject of "Board Tax Documents." The message matched two policies—one that matched the untrusted and enforced (moved) policy and another that looked to be a spoofing attempt. This message is shown in Figure 16.

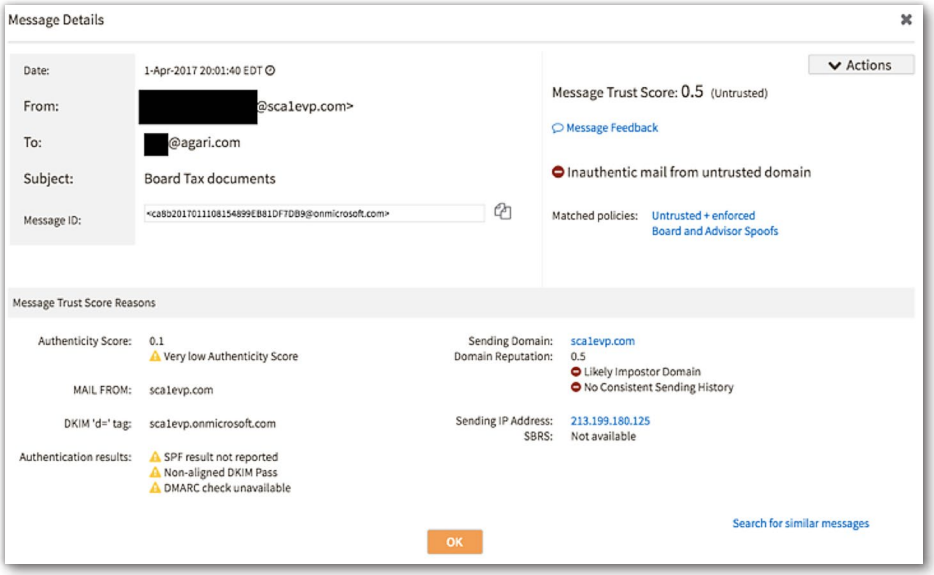

Figure 16. A Malicious Message Caught by Agari

As can be seen, Agari was able to detect several "tells" that this message is bogus:

- The message is from a likely impostor domain (sca1evp.com).
- The message has not been seen in legitimate messaging before.
- The message has a low authenticity score.

As with any other message, we could immediately build a policy based on any of its attributes (domain, sender, IP, content, etc.) or simply notify security teams and responders.

Searching for other suspicious attributes of messages is simple in the Agari Messages Dashboard, as seen in Figure 17, where we looked specifically for other "impostor domains" noted in the past month.

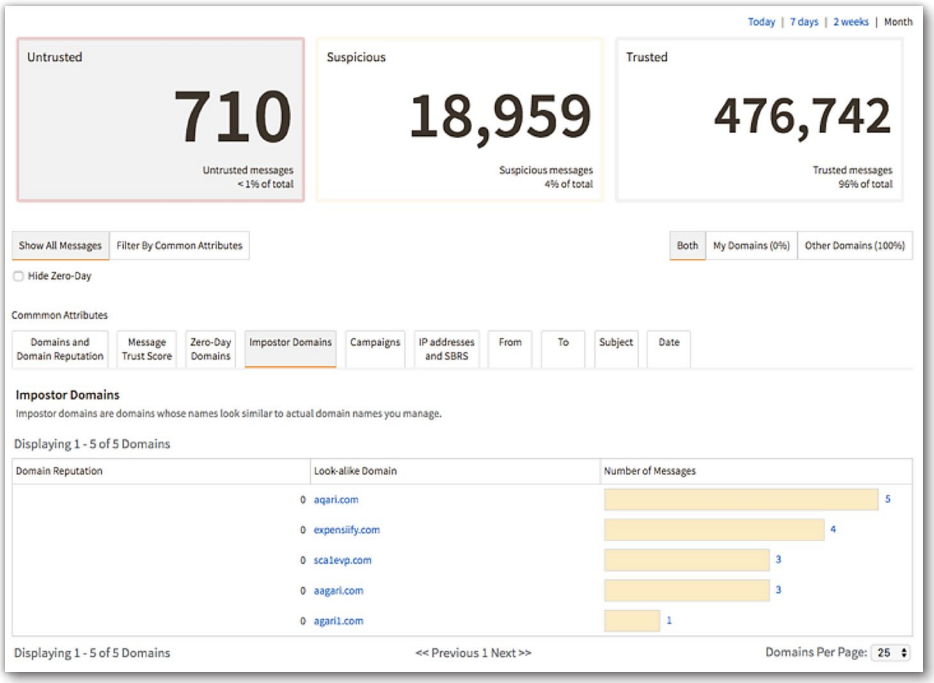

Figure 17. Searching for Additional "Impostor Domains"

With this feature, we were able to see additional domains that were trying to spoof the Agari domain and Agari department domains, along with the number of messages coming from those domains.

With this information, we conducted similar "attribute" searches to look for untrusted emails to specific recipients, such as the executive in the previous example. With highly targeted, malicious campaigns rising against senior executives and influential users who might have access to sensitive assets and data, building highly specific prevention policies and alert configurations that focus on suspicious content that targets these users makes a lot of sense.

Finally, we reviewed some of the reporting capabilities in the Agari console by selecting Manage/Reports. The initial page we landed on shows the category breakdown of messages matching defined policies, both seen and moved. By clicking the simple graphs presented on this screen, we can see more detail about any of the policy periods. Figure 18 shows events of the "untrusted+enforced" policy over two weeks at Agari.

TAKEAWAY:

By tracking the changes in messages detected and possibly moved within a time range, we can fine-tune policies to be more accurate and improve detection capabilities. Ultimately, this helps reduce risk over time.

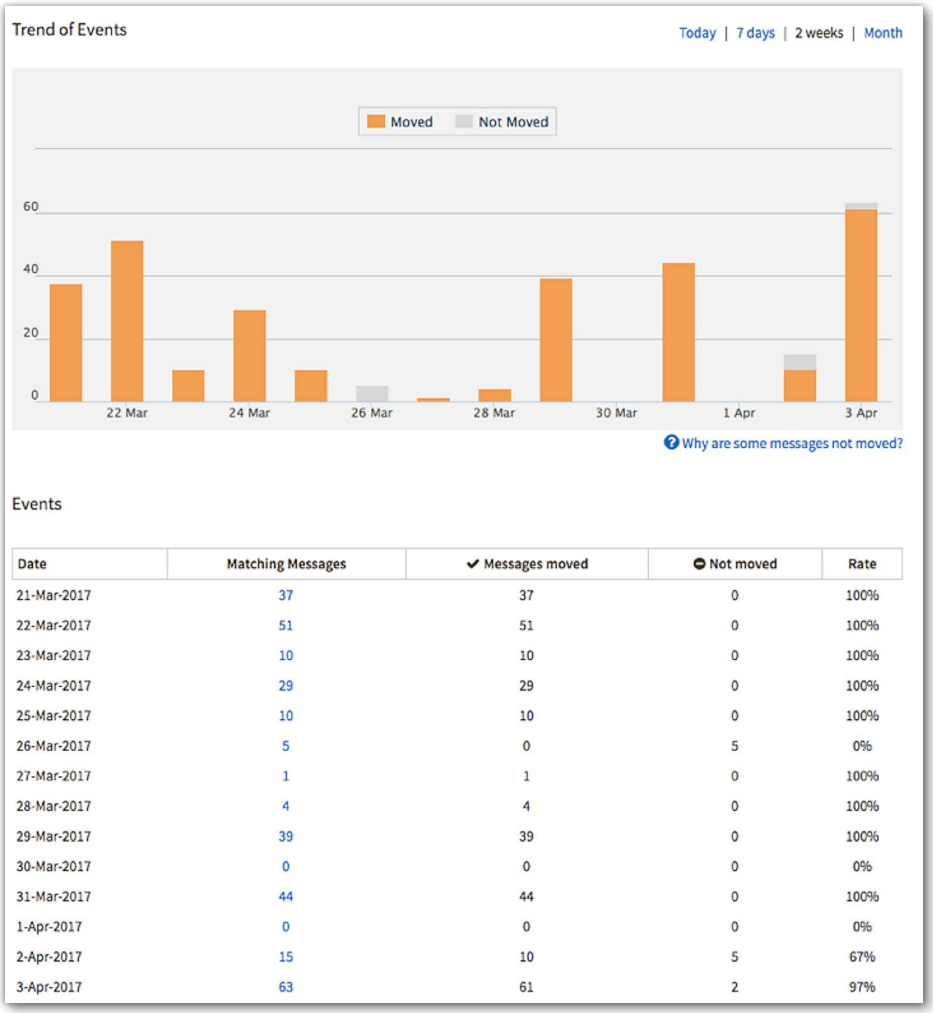

Events

Date	Matching Messages	✔ Messages moved	⛔ Not moved	Rate
21-Mar-2017	37	37	0	100%
22-Mar-2017	51	51	0	100%
23-Mar-2017	10	10	0	100%
24-Mar-2017	29	29	0	100%
25-Mar-2017	10	10	0	100%
26-Mar-2017	5	0	5	0%
27-Mar-2017	1	1	0	100%
28-Mar-2017	4	4	0	100%
29-Mar-2017	39	39	0	100%
30-Mar-2017	0	0	0	0%
31-Mar-2017	44	44	0	100%
1-Apr-2017	0	0	0	0%
2-Apr-2017	15	10	5	67%
3-Apr-2017	63	61	2	97%

Figure 18. Policy Reporting over a Time Period

Any of the "Matching Messages" counts can be clicked to take you to a list of the actual messages detected for that day.

Conclusion

Email is one of the top threat vectors we face today, and it deserves more attention from security teams. We need to implement effective email security controls that can be readily managed and provide the level of visibility into malicious email campaigns to effectively block them before they lead to incidents and breaches.

After reviewing Agari's Advanced Threat Protection platform, SANS found the product to be simple to use and manage. Policies were easy to create, and Agari's advanced analytics for email trust accurately identified malicious and suspicious emails. The console was simple to navigate, with strong risk visualization details and drilldown capabilities. For enterprises looking for a solution that can help significantly reduce the onslaught of malicious email to our users today, Agari Advanced Threat Protection seems to be a solid option.

About the Author

Dave Shackleford, a SANS analyst, instructor, course author, GIAC technical director and member of the board of directors for the SANS Technology Institute, is the founder and principal consultant with Voodoo Security. He has consulted with hundreds of organizations in the areas of security, regulatory compliance, and network architecture and engineering. A VMware vExpert, Dave has extensive experience designing and configuring secure virtualized infrastructures. He previously worked as chief security officer for Configuresoft and CTO for the Center for Internet Security. Dave currently helps lead the Atlanta chapter of the Cloud Security Alliance.

Sponsor

SANS would like to thank this paper's sponsor:

agari
by HelpSystems