A Layered Approach to Defending Against List-Linking Email Bombs

Cristina Houle
Forcepoint Security Labs
Forcepoint
San Diego, USA
choule@forcepoint.com

Ruchika Pandey
Forcepoint Security Labs
Forcepoint
San Diego, USA
rpandey@forcepoint.com

***Abstract*—Email bombing is a form of Denial of Service (DoS) attack that consists of sending huge volumes of email to one or more email addresses to overflow the mailbox or overwhelm the server where the mailbox is hosted. While this type of attack is not new, we have seen renewed Distributed Denial of Service (DDoS) attacks by email in the past few years via list-linking attacks, where the victim's email address is subscribed to thousands or even tens of thousands of mailing lists. We describe some of the measures that list owners and email service providers have suggested to mitigate such attacks. We then present a case study of a real-world attack, investigating whether characteristics of attack behavior and email attributes surface a workable hypothesis of early detection paradigms. We test our hypothesis on a dataset of hundreds of millions of emails, representing three months of data from 3,000 medium and large organizations. Our technique helped us detect three previously unknown list-linking attacks in the dataset. We determined that anomalous bursts signify meaningful patterns at the user level, but cannot be extrapolated effectively when analyzing bulk enterprise volume data. Effective spam identification in this case needs to consider the unique nature of the messages in the attack, which lends to their exhibiting linguistic similarities, combined with their temporal proximity. We recommend a layered approach to detection and throttling through per-user volume and time-based methodologies paired with phrasal pattern recognition.**

Keywords— Anomaly detection, Computer crime, Cyberattack, Distributed denial-of-service attack, Internet security, Unsolicited electronic mail

I. INTRODUCTION

Email spam is an intrusive, pervasive, and resource-draining distraction that impacts entities at every level, from individuals to business to government. Spam represents the bulk of email sent daily [1], and even as spam-filtering efforts to detect and deflect evolve, so too do the evasive attempts of spammers. However, while email spam is relentless and taxes the bandwidth of recipients, inboxes, spam filters, and servers, the aim – even by phishers and malware senders – is generally to be read.

This is not the case with an email bomb, which floods its recipients with an unusually high volume of email over a period of time, anywhere from hundreds per hour to hundreds per minute [2]. When email is weaponized in this fashion, it operates as a Distributed Denial of Service (DDoS) attack on the individual user or email server, spam filter, or other resource, crippling communication and potentially derailing basic functions for the entire entity. The end goal here is not for the recipient to receive the message, but to render the mailbox – or entire mail system – unusable.

In the early days of email as a communications tool for individuals and businesses, it did not take long for agents of chaos to adapt and exploit it as an attack vector. Disruptive email messaging by volume was quickly recognized and used by individuals and organizations as an instrument of activism, retaliation, avarice, and more.

Initially, targets were typically governmental. One of the first documented email bombs was launched against the Langley Air Force Base (AFB) email infrastructure in 1997 [3]. In 1998, the Tamil Tigers swamped Sri Lankan embassies with 800 emails per day over a 2-week period, with messages that claimed responsibility and stated their objective in the message body: "We are the Internet Black Tigers and we're doing this to disrupt your communications [4]."

Twenty years later, email bombs are more sophisticated, prevalent, devastating, and varied in their targets. Today's email bombs are increasingly list-linking email bombs, which consist of high volumes of confirmation messages from forums and message boards, newsletter signups, retail mailing lists, and other everyday communications familiar to most users. ("Confirmation messages" are used here to describe acknowledgements of joining or subscribing; congratulatory and similar exclamatory statements; coupons and other inducements to purchase; and requests for further action.)

During an email DDoS attack, these messages typically confound and overwhelm most spam filters, not only due to the sheer number of emails received, but also because they are legitimate, safe messages from benign sites functioning in the commercial and community sphere. These sites, some of which even use recognized marketers to facilitate their outreach, thereby become an unwitting conduit for a sustained attack. Such attacks no longer need to make bold, overt, declarative statements of malicious intent in the torrent of messages. Instead, as Marshall McLuhan observed, the medium is the message [5].

978-1-5386-4922-0/18/$31.00 ©2018 IEEE

Our work contributes a new approach to dealing with list-linking attacks while still allowing regular email, including most marketing messages, to be delivered. We first discuss previous research and recommendations, and address the challenges to those approaches. We then examine a list-linking attack in detail and build a hypothesis around identifying attacks by taking advantage of the shared infrastructural components of the Internet that these attacks leverage. We explain our methodology and our conclusions regarding why a User and Entity Behavioral Analytics (UEBA) based approach is needed. We highlight the phrasal patterns present in the traffic, and suggest that phrasal pattern recognition can help confirm the nature of the attack and throttle the volume by identifying email messages that should be characterized as spam in this case, but not necessarily at other times or for other users.

II. List-Linking Email Bomb

The locus of the list-linking email bomb is distributed, emanating from sources located around the globe and encompassing an array of languages, from Albanian to Turkish, Catalan to Greek. The advantage of distributed attacks, of course, is a multiplicity of origin that camouflages intent and obscures the true attacker while effectively overwhelming human and machine resources. The advantage of the simple web form – a contact page, a forum signup, a newsletter or store subscription form – as a vehicle for attack is its ubiquitous universality and structural similarity. Built using Joomla, Wix, Wordpress, or any number of similar aids, each nondescript and unvalidated form, on sites old and new, fulfills a unique function as an access point to potentially paralyzing disruption in the aggregate.

A list-linking email bomb leverages the commerce, communication, and community-building tools and resources used by small businesses, charities, and communities – groups that lack the capital; knowledge and technical resources; and motivation to implement defensive measures such as Completely Automated Public Turing Test To Tell Computers and Humans Apart (CAPTCHA) and other inhibitors to abuse and manipulation.

However, it is this exploitation of infrastructural underpinnings that yields a possible solution in the semantic uniformity of confirmation messaging, despite minor linguistic dissimilarities. Combined with anomalous time and volume signifiers in an examination of UEBA, and paired with similarities across sender addresses at the mailbox and domain level, recognizable patterns of attack seem to emerge early enough to be thwarted in the initial stages, possibly as soon as the first or second hour in what is typically a 12- to 24-hour siege. Early detection in such attacks is critical to an effective defensive strategy, as it can help prevent systems from being overwhelmed and knocked offline. These same signifiers can be used to identify attack abatement as well, by detecting the sharply decreasing volume as it returns to pre-attack levels.

In the absence of this more nuanced approach, the options for organizations, email providers, and spam filters are limited to more draconian measures such as blacklisting large mailing list providers [2], disabling targeted mailboxes [6], or aggressively dropping all email marketing messages. Additionally, as these attacks become commoditized and therefore easier to execute or outsource [7], they are increasingly targeted at individual users where their volume does not necessarily exhibit as a statistically significant anomaly at the aggregate organizational level.

III. Previous Research and Recommendations

A. *Achilles' Heels: Web Forms and Subscription Pages*

List-linking email bombs co-opt the communication and recruitment tools of the web through several means: scripts and agents; Dark Web mercenaries; and sites such as research-friendly mailbait.info (and its less noble counterparts) that have aggregated vulnerable forms and consolidated scripting tools into a single, simple action that requires no registration and little or no payment to launch.

As detailed by Jakobsson and Menczer in 2003 [8], this ease of access via unsecured subscription signups highlights the specific and alluring advantages to using a standard communication channel such as email to effect a system takedown: It is a low- or no-cost venture that does not involve intrusion or infiltration, meaning such attacks are untraceable. These factors make the email bomb a high-reward, low-risk disruptive activity that will likely cause it to remain a favorite attack vector [2].

B. *Defending Against Unprotected Signups*

After several notable incidents of email bombs in the late 1990s, including episodes involving high-profile, nation-state targets as diverse as the Sri Lankan diplomatic corps[4], the US Air Force [3], and the Yugoslavian government [4], research efforts into email bombs and web form vulnerability produced recommendations for several possible solutions to be integrated during the signup process.

CAPTCHA was one of the earliest proposed solutions to verify human engagement [8]. It has since been improved upon with RECAPTCHA and other modifications or behavioral variants, such as image recognition CAPTCHA. While arguably boasting the highest adoption rate of any proposed effort to hinder machine agents at the time and point of engagement [6], CAPTCHA does not enjoy a 100% adoption rate and therefore, from the recipient's perspective, cannot be considered entirely effective.

Other proposals included Confirmed Opt-In (COI) efforts [8]. These began as confirmation messages requiring registration validation through active link clicking, verifying the recipient's willing participation. Subsequent improvements to COI methodology involved mailto links generated at registration.

However, the first and still most prevalent form of COI, confirmation messaging, merely compounds the problem faced by victims of a DDoS email bomb attack, by generating confirmation requests that swamp the recipient's inbox and increase server load issues. The advantage to COI lies over the long-term, which is little comfort to an individual experiencing real-time distress and lost emails. When the attack is over and

the email flow subsides to a trickle, the recipient benefits from the fact that active confirmation is required to receive more email from those particular sources.

As outlined by Jakobsson and Menczer [8], the second approach using dynamic mailto links can shut down simple scripts, but also has its own set of secondary steps that could be necessary to be completely effective in evading savvy scripts, agents, and bots. In the end, no matter how effective the multi-step validation process of a self-generated mailto link or CAPTCHA form might be, the overriding issue is that the barrier to adoption of such a method is very high for most webmasters and administrators of small business and single-interest sites such as blogs, forums, and fan sites.

Since being referenced in Jakobsson and Menczer's 2003 paper [8], none of these methods has experienced full adoption [9]. Plugins and other add-on modules developed by some of the more popular web software platforms are not the default, and many users fail to implement such measures due to lack of awareness or technical knowhow. To facilitate greater compliance, open source software developers would need to prioritize and develop such tools, and make them easier to deploy. Additionally, many sites and services desire to have an unconstrained subscriber acquisition process due to the financial and other inducements a large subscriber database offers. One example of this commercially-influenced resistance can be seen in MailChimp's 2017 decision to revert to Single Opt-In (SOI) as their default [9]. Moreover, older sites with these web forms continue to exist, often unmaintained and never upgraded. They represent a meaningful component of the threat presented by form-initiated DDoS email attacks.

C. Shift from Signup to Sendout

Anecdotal studies of recent email bomb attacks [10] and subsequent recommendations for attack detection and disruption resulted in an Internet Engineering Task Force (IETF) proposal for a new header, the Form-Sub header [11].

The Form-Sub header, proposed in 2017, suggests integrating the subscriber's Internet Protocol (IP) address in the header, obfuscating part of the IP to protect personally identifying information (PII) [12]. This header could help the recipient's mail server to identify similarities in the IP-based origins of an email deluge and help to eliminate chokepoint vulnerabilities, but its efficacy seems predicated on the IP of the subscriber being identical in every email. This obviates the header's utility in sophisticated, concerted attacks launched by a network of bad actors, as might be experienced by users victimized by nation-state actors or coordinated hacking networks.

The proposal for the Form-Sub header is under consideration and remains in draft form, but our position is that a significant impediment to its success, from the standpoint of an email bomb recipient, relies on its adoption by all sites with forms. While larger players in the email marketing sphere are open to such mitigation attempts (MailChimp is one vendor that has signaled its intent to implement [6], even while undermining its potential efficacy by returning to SOI as their default [9]), the probability that sites with small footprints are unlikely to incorporate such strategies represents a significant challenge. These smaller sites often lack the economic and technical wherewithal to implement such proactive measures; are often unaware such measures exist, even if offered as plugins by common platforms; or intentionally try to avoid limiting membership growth with validation efforts such as CAPTCHA or COI, etc.

IV. Case Study

Our case study examines an incident involving a user account that was known to be the victim of a list-linking email bomb. In addition to an abrupt increase in email volume, the data indicate that anomalies in other attributes such as origin and language can also be assessed as attack indicators. Of these, the geographic distribution of email senders, as well as the language of the email subjects, were evaluated and determined to be the most salient.

The theories developed from that analysis were then tested on our full dataset of 3,000 enterprise accounts. This helped us identify three additional user mailboxes where existing spam detection could have been augmented by earlier identification as a list-linking email bomb. We looked for volume-based anomalies where an anomaly was indicated by an increase in number of emails received per day equal to 10 standard deviations from the mean over the previous 7 days. We used a combination of linguistic pattern analysis using key phrases along with an anomalous jump in the volume of emails to identify accounts being targeted by email bombs.

Our case study underscores the need for (and potential efficacy of) a layered approach to detection and throttling. The initial layer leverages the identification of user-based volume anomalies within sharply delineated time periods as a potentially meaningful event signifier. This indicator suggests a need for deeper analysis, achieved through the incorporation of phrasal pattern analysis to discern the materialization of a list-linking email bomb. Additional supporting indicators such as anomalies in other metadata like geographic origin, subject language, and repetitive sender features can also help validate the detection of a list-linking attack, providing the basis for a method to identify and drop emails that might otherwise overwhelm and incapacitate the user's mail system.

A. Time and Volume

As seen in Fig. 1, list-linking email bombs are marked by the rapid onset and sustained delivery of confirmation messages over a period of time. In most cases, the uptick in volume is sharp and sudden relative to normal day-to-day use and hourly averages, reaching the volume apex within the first few hours. There is no slow incrementation over time; it is acute.

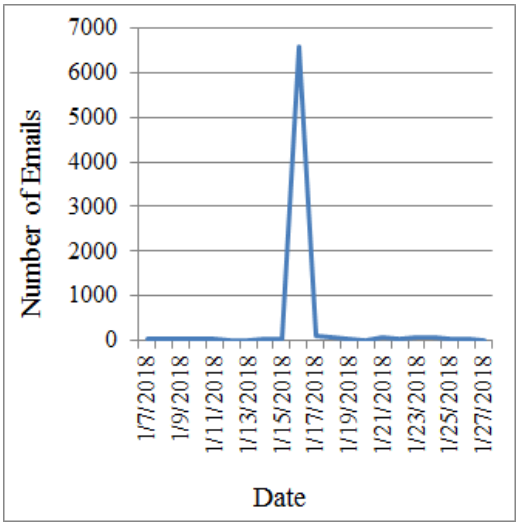

Fig. 1. Daily Email Volume Over 3 Weeks

As Table I shows, once the zenith is reached, the threshold remains high. Similar to the onset, the conclusion of the attack does not taper, but rather plunges precipitously. The data show that the bulk of the attack consisted of COI emails, which caused a spike in volume of email, and that the sustained intensity spanned a period of hours rather than days.

TABLE I. HOURLY EMAIL VOLUME

Time	Number of Emails
12:00:00 AM	1
1:00:00 AM	93
2:00:00 AM	763
3:00:00 AM	828
4:00:00 AM	737
5:00:00 AM	730
6:00:00 AM	729
7:00:00 AM	624
8:00:00 AM	511
9:00:00 AM	543
10:00:00 AM	446
11:00:00 AM	466
12:00:00 PM	65
1:00:00 PM	5
2:00:00 PM	5

In our example, shown in Table II, a single mailbox was targeted and received over 6,500 messages within a 12-hour period. Once the deluge began, the end user did not receive a single legitimate message for more than 9 hours. Table I shows the number of messages jumped from fewer than 5 to close to 100 in the first hour of the attack. Within the first 10 minutes, the user received 30 confirmation messages. By the third hour, the number of confirmation messages peaked at over 800. The primary thrust of the attack concluded in the 12th hour, with the number of emails dropping sharply to just 65. After that, the flow of emails trickled in at a low rate per hour.

TABLE II. DAILY EMAIL VOLUME

Date	Number of Emails
1/12/2018	12
1/13/2018	5
1/14/2018	27
1/15/2018	27
1/16/2018	6571
1/17/2018	100
1/18/2018	60
1/19/2018	25
1/20/2018	14
1/21/2018	78

B. Geography

Fig. 2 shows the volume of email by country of the sending Mail Transfer Agent (MTA) on the days surrounding the

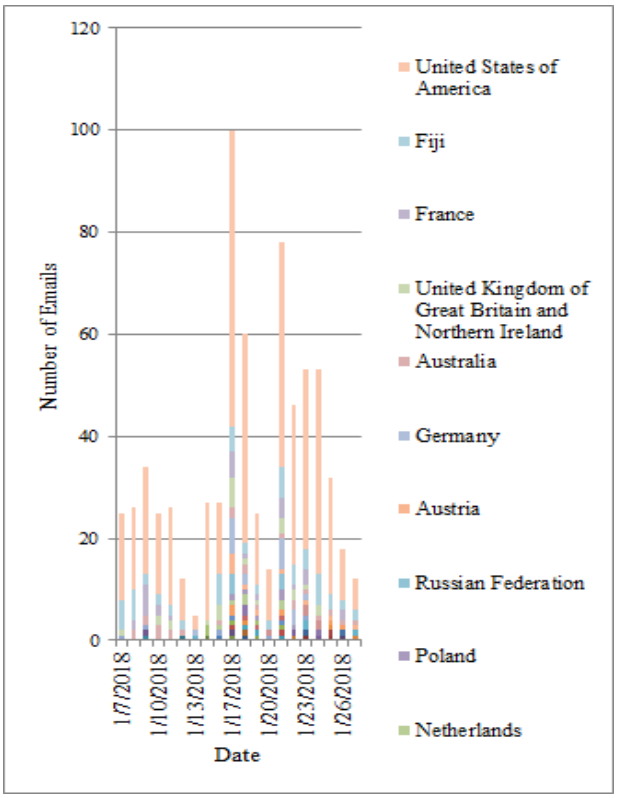

Fig. 2. Sender MTA Countries in Attack

attack. Preceding the attack, emails were received from senders in fewer than five countries. In contrast, senders from 87 countries were involved in the attack, as seen in Fig. 3. The lingering effects of SOI artifacts from an email bomb can be seen in the ongoing arrival of foreign-language spam after the main attack was concluded. Where COI spam trails off, even after multiple attempts at confirmation, SOI spam requires the user to actively engage with the email – for instance, by blacklisting the sender or unsubscribing – in order to stop receiving mail.

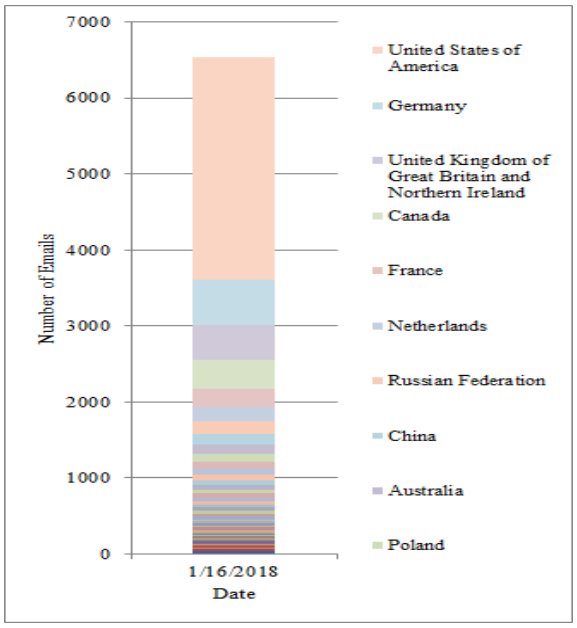

Fig. 3. Email Volume by Sender MTA Countries

C. Language

We can see from Fig. 4 that emails were entirely in English before the attack. The attack introduced emails in various languages, and those with the highest representations are listed in Fig. 4 and Fig. 5. Even after the bulk of the attack was over, we saw some of the languages continue to appear in the mailbox, mostly the residual effect of SOI subscriptions (coupons, sales announcements, etc.).

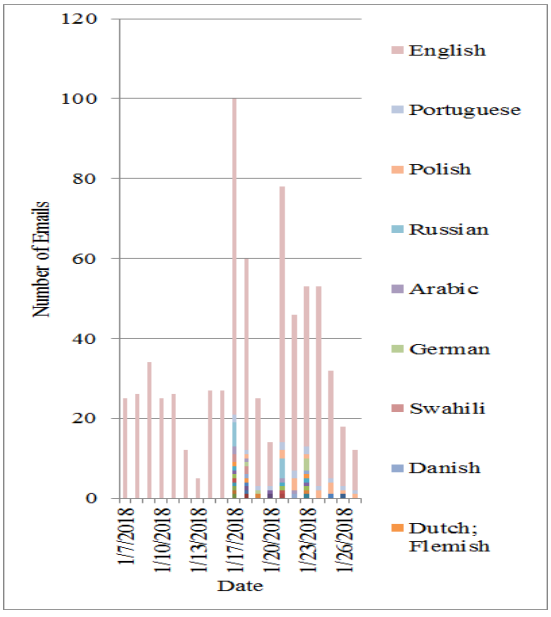

Fig. 4. Linguistic Variation by Subject, Pre- and Post-Attack

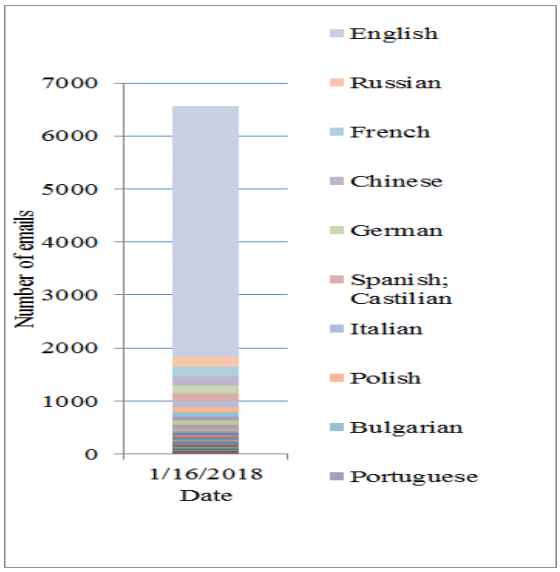

Fig. 5. Linguistic Variation by Subject in the Attack

D. Linguistic Patterns

Based on our case study we determined that robust anomaly detection and a thorough understanding of UEBA can be used to identify and foil such attacks, deflecting attempts to incapacitate a user's inbox or paralyze the network.

While English represents about 74% of all the emails received, as seen in Fig. 6, one cannot dismiss the messages received in roughly 33 other languages. As email bombs avail themselves of existing, unguarded, and generally unmodified subscription modules that support a broad scope of activities from ecommerce to single-issue communities, the language

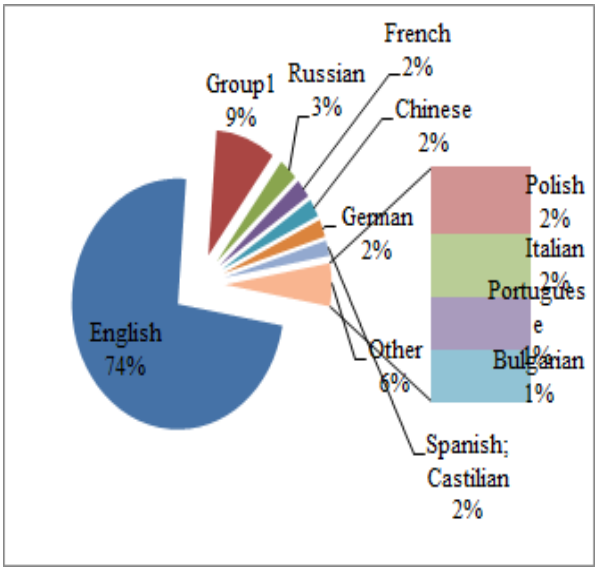

Fig. 6. Top Languages by Email Volume

that supports these exchanges is significantly standardized. Despite some linguistic variances in messaging, the ability to recognize semantic patterns can therefore contribute meaningfully to rapid detection.

Even when fragmented into components, many languages are represented in significant enough volume – and use expected, recognizable semantic patterns – that discernment and discardment should be possible, especially when combined with volume and time-based signifiers.

For instance, in the course of the 12-hour attack in this case, nearly half – 44.8% – of the confirmation messaging subject lines were in English and began with "Account details," with 7% concluding that subject line with "(pending admin approval)." Other commonly occurring phrases are shown in Table III.

An additional 9% of the overall volume began with foreign variants of this same phrase, with "Kontoinformationen," "Détails du compte," and "Szczegóły konta" accounting for some of those numbers.

TABLE III. EMAIL SUBJECT OCCURRENCES

Number of emails	Subject starts	Subject contains
2948		Account details
740	[	
734	Welcome	
583	Welcome to	
534		pending
533		approval
441		regist
332		Forum
302		registration
302		Your username
302		username and password
217		activat

Other significant numbers include 11% of subject lines starting with a "[," followed by the name of the forum/site. Overall, 4% (302) began with "[" followed by the forum name, and contained "username and password" such as "[Newland Travel] Your username and password." Another 11% began with the word "Welcome," as seen in "Welcome to IJTTE Forum."

E. Repetitive Sender Features

In keeping with the distributed nature of the attack, we see that the emails were sent from a wide variety of IP addresses. On the day of the attack, 6,555 emails were sent by 4,047 unique IP addresses. Sender Autonomous System Numbers (ASNs) show a pattern in line with the theory that these emails were sent from a variety of legitimate web properties all over the world. Table IV shows the top 10 ASNs of email senders by volume.

TABLE IV. TOP 10 ASNS BY EMAIL VOLUME

ASN	Number of Emails	Name
AS26496	653	GODADDY
AS8560	442	ONEANDONE-AS Brauerstrasse 48, DE
AS36483	237	GOSSAMERTHREADS - Gossamer Threads Inc., CA
AS26347	216	DREAMHOST-AS - New Dream Network, LLC, US
AS16276	196	OVH, FR
AS24940	190	HETZNER-AS, DE
AS29873	120	BIZLAND-SD - The Endurance International Group, Inc., US
AS20738	118	AS20738, GB
AS46606	112	UNIFIEDLAYER-AS-1 - Unified Layer, US
AS32475	87	SINGLEHOP-LLC - SingleHop, Inc., US

The volume of confirmation messages from diverse sites means that sender email addresses are unique, but the appearance of certain features of the sender's email address – before and after the @ – can also act as indicators that help to identify the onset of an email bomb DDoS attack. These features, which are not typical of emails received by most end users, especially in high volume, can assist in confirming the attack context of certain phrasal fragments.

For example, a meaningful number of emails with a subject line beginning with "Welcome to" or "Account details" may signify a representative part of an attack and can be identified as anomalous behavior through high frequency occurrence within a limited time span. The appearance of address components such as mailbox names like "info@" or domain name like ".secureserver.net," especially in statistically significant numbers, can validate this interpretation. Fig. 7.

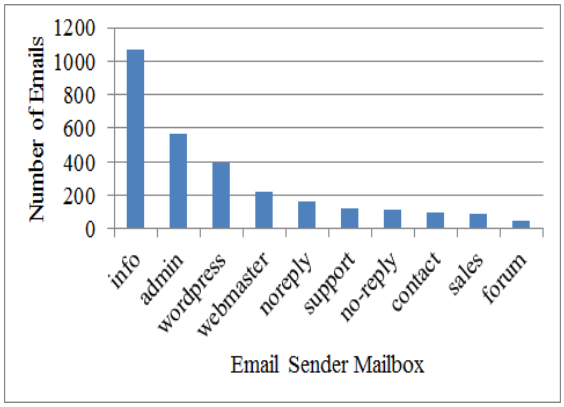

Fig. 7. Top 10 Email Sender Mailboxes by Volume

shows the top mailbox names involved in the attack.

Additionally, their presence can assist in preventing False Positives (FPs) arising from volume, duration, and phrasal pattern recognition. The presence of 4 emails containing .secureserver.net in the from address as well as a subject line

starting with “Welcome to” in the first 10 minutes of a DDoS email attack, as exemplified by our dataset, can offer both a redundant signal of anomalous behavior and confirm the substantive context is that of confirmation messaging.

In our case study, for instance, the first 30 messages in the 12-hour attack, received in a 10-minute window, contained 4 messages from “admin@” and 3 from “nobody@.” While this is already highly unusual, when combined with signifiers such as the subjects of these particular messages, this combination of elements provides an additional layer of detection and therefore a higher degree of accuracy.

V. Analysis and Results

A. Enterprise Anomaly Detection Using Aggregate Volume

We started our study with one known email bombing attack. To determine how attack detection could be automated effectively, we analyzed three months of email data for 3,000 medium and large enterprises, ranging from community colleges with less than 200 users to larger multi-nationals with over 10,000 users. Analysis looking for volume anomalies at the aggregate level surfaced 12 enterprise accounts for further processing. We then conducted a detailed analysis of per-user volume, headers, subjects, and other metadata for these 12 enterprise accounts encompassing 34,604 users and 10,689,761 emails.

1) Methodology

The first step was to calculate the mean of the total number of emails received per day over the preceding 7 days and the preceding 30 days for each of the 3,000 enterprises, and review sudden spikes in email volume. We chose 7 days and 30 days to normalize over weekly and monthly variations in work patterns. This method yielded 12 candidates, which received an average of at least 200 emails per day, and where there was a single day jump in total number of emails received equal to at least 100 standard deviations from the mean. Next, we analyzed these accounts in detail to see if they were the target of email bombs. We calculated the 7- and 30-day mean of emails received by each user, and looked for days where the number of emails received on a given day were at least 10 standard deviations from the mean.

We found some minor increases in volume – up to 10 times the standard deviation - caused by responses to meeting invites, email receipts, farewell emails, etc. In most, but not all, cases the senders and recipients were on the same email domain. While contributing to a greater volume of email, the closed-loop insularity of intradomain communication acts as a buffer against misinterpretation of these emails as part of an attack.

Larger spikes – sometimes a few hundred times the standard deviation – were caused by automated notifications from business applications like Sharepoint, repeated process logs, and status and error updates from various hosted domains. Another set of email accounts yielding large single-day spikes – over 1,000 standard deviations from the mean – were spam and security notification accounts. This group included large numbers of emails from the same senders, and had either identical or partially matching subjects. These patterns help eliminate these users from being identified as potential victims of email bombing attacks.

From this analysis of anomalies using high-volume data, we determined that the presence of behavioral anomalies at the user level, even when statistically significant and potentially devastating as with a targeted email bomb, is effectively obscured by the higher-volume patterns experienced by the organization.

B. User-Based Anomaly Detection

The next step was to conduct the same analysis looking for per-user statistical anomalies. This helped us identify 32 individual email accounts of interest that had not been surfaced when using enterprise-level anomaly detection as an initial filter. Further review of senders and subjects, including linguistic analysis looking for phrases such as "Account details" in the subjects helped separate these data into two types of mailboxes. Of these, 29 mailboxes received various types of spam or business process notification. Three accounts were hit by list-linking email bombs.

The first mailbox identified by our methodology belonged to a user at a community college. As can be seen in Fig. 8, this user went from receiving no email in the week prior to the attack to a peak of 3,262 emails in a single hour, and then back to receiving a few emails per week.

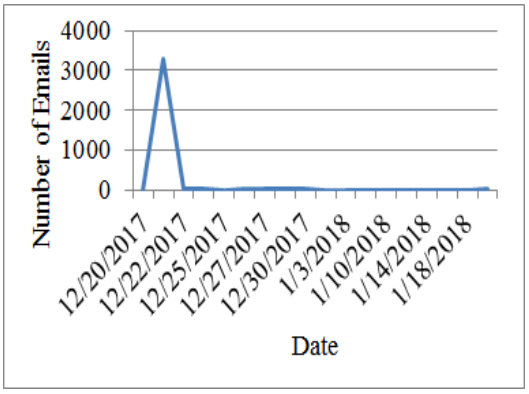

Fig. 8. Daily Email Volume for First User

The second mailbox belonged to a user at a large multinational. The attack on this mailbox, as shown by Fig. 9, manifested similar characteristics, with 3,806 emails received in one hour. This represented a significant change in volume from the 20 to 30 emails received per day on previous days. The volume of emails remained relatively high for about 4 days before settling to a new normal of around 60 per day.

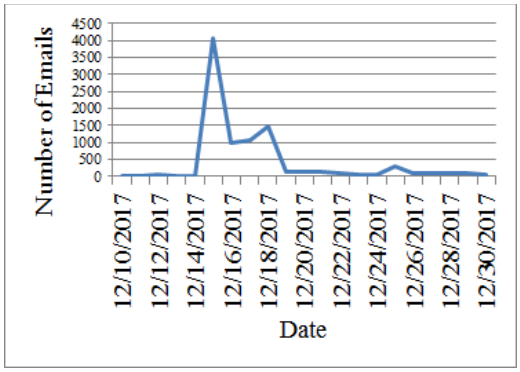

Fig. 9. Daily Email Volume for Second User

The last mailbox similarly went from receiving less than 20 emails per day to 1,381 emails in a single hour, as can be seen in Fig. 10. The attack lasted only a few hours, after which the email volume returned to previous levels.

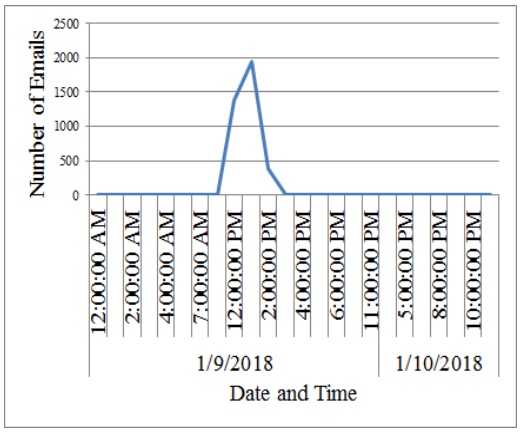

Fig. 10. Daily Email Volume for Third User

VI. Practical Application

An abrupt change in email volume can be determined by calculating a rolling average of hourly email volume per user. A sudden, sharp deviation in volume, coupled with phrasal pattern recognition, would indicate the potential onset of an attack. Once an attack is suspected, a combination of heuristics – semantic patterns, language of the subject, and the geography of senders – can be used to confirm the onset of the attack and move to a heightened defensive posture. In this state, suspected spam emails would be dropped more aggressively to prevent the mailbox from being flooded. Monitoring these same elements can be used to detect the end of the attack as well, and inform the decision to drop down to a normal operating posture. For instance, the daily or hourly volume of emails dropping to three standard deviations from the mean for the period prior to the onset of the attack and staying at that level for 3 to 5 days could be used as an indicator of the end of the attack. Dropping back down to a normal defensive posture is important to prevent false positives in normal ongoing mailing list traffic.

VII. Conclusion

There is concern among security industry experts that email bombing, already proliferating, will continue to increase in frequency and may come to represent a major tool in the disruptor's arsenal [2],[7]. The cost benefit analysis (CBA) of the email bomb as a weapon to be wielded against governments, businesses, the press, and individuals has likely already been calculated and found favorable. Untraceable, effective, and inexpensive, the email bomb is being automated and disseminated with such efficiency that Spamhaus referred to the monetized threat as Mail-bombing as a Service (MaaS) [2].

To date, most proposed solutions to the growing danger posed by email bombs as highly effective service disruptors have revolved around defensive measures initiated at the point of origin, such as CAPTCHA or COI. Less attention, however, has been focused on defensive measures at the email bomb's point of impact.

We believe that in cases where the email bomb is of low or moderate volume targeting an individual user, or where the onset exhibits a manageable increase in the nascent stages, a combination of user email behavior profiling and anomaly detection can be used to identify the start of an email bombing attack very quickly. Early detection can help to maintain the functionality of a user's individual mailbox, supporting the ongoing processing of expected messages while initiating a heightened defensive posture to mitigate against mailbox and server overload. Identification in the preliminary stages of an email DDoS attack can ensure that more drastic measures, such as IP blocking or the disabling of mailboxes or servers, are unnecessary. Additionally, defensive strategies informed by an understanding of UEBA can contribute to a similar observation of abatement detection, providing data integral to determining the end of the attack as well.

Acknowledgment

The authors thank Forcepoint Security Labs team members for their support, encouragement, peer reviews, and feedback.

References

[1] Forcepoint, "Forcepoint 2016 Global Threat Report," Forcepoint, Apr. 2016. [Online]. Available: https://www.forcepoint.com/resources/reports/forcepoint-2016-global-threat-report

[2] L. Maltice, "Subscription Bombing: COI, CAPTCHA, and the Next Generation of Mail Bombs," Spamhaus, Sep. 12, 2016. [Online]. Available: https://www.spamhaus.org/news/article/734/subscription-bombing-coi-captcha-and-the-next-generation-of-mail-bombs

[3] T. Bass, A. Freyre, "E-Mail Bombs and Countermeasures: Cyber Attacks on Availability and Brand Integrity," IEEE Network Magazine, Vol. 12, No. 2, pp. 10-17, March/April 1998. [Online]. Available: https://www.thecepblog.com/2016/07/23/papers/pdf/ieee-network-email-bombs.pdf

[4] D. E. Denning, "Activism, Hacktivism, and Cyberterrorism: The Internet as a Tool for Influencing Foreign Policy", Global Problem

Solving Information Technology and Tools, December 10, 1999. [Online]. Available: https://nautilus.org/global-problem-solving/activism-hacktivism-and-cyberterrorism-the-internet-as-a-tool-for-influencing-foreign-policy-2/

[5] M. McLuhan, "The Medium is the Message," in Understanding Media: The Extensions of Man, New York, NY, USA: McGraw Hill, 1964, pp. 1. [Online]. Available: http://web.mit.edu/allanmc/www/mcluhan.mediummessage.pdf

[6] J. Angwin, "How Journalists Fought Back against Crippling Email Bombs," Wired, Nov. 2017. [Online]. Available: https://www.wired.com/story/how-journalists-fought-back-against-crippling-email-bombs/

[7] D. Pickett, "Email Bombs Increasing in Frequency", appriver, Jan. 10, 2018. [Online]. Available: https://blog.appriver.com/2018/01/email-bomb-attacks-increasing-in-frequency/

[8] M. Jakobsson, F. Menczer, "Untraceable Email Cluster Bombs: On Agent-Based Distributed Denial of Service," CoRR, cs.CY/0305042,pp. 1-11, May 2003. Doi: arXiv:cs/0305042v1 [cs.CY], [Online]. Available: https://arxiv.org/pdf/cs/0305042v1.pdf

[9] MailChimp, "Why Single Opt-In? And an Update for Our EU Customers," Oct. 30, 2017. [Online]. Available: https://blog.mailchimp.com/why-single-opt-in-and-an-update-for-our-eu-customers/

[10] B. Krebs, "Massive Email Bombs Target .Gov Addresses," KrebsOnSecurity, Aug. 16, 2016[Online]. Available: https://krebsonsecurity.com/2016/08/massive-email-bombs-target-gov-addresses/

[11] Messaging Malware Mobile Anti-Abuse Working Group (M3AAWG), "M3AAWG: M3AAWG Recommendation on Web Form Signup Attacks," 2017. [Online]. Available: http://www.m3aawg.org/WebFormAttacks

[12] J. Levine, IETF Network Working Group, "A Message Header to Identify Subscription Form Mail," 2017. [Online]. Available: https://datatracker.ietf.org/doc/draft-levine-mailbomb-header/?include_text=1