Subscription Bombing: Email Under Attack

MATHIAS PAYER, EPFL, Switzerland
CANDID WÜEST, xorlab, Switzerland

Email subscription bombing, also known as subscription flooding, is an emerging attack vector in which an attacker subscribes the victim to thousands of mailing lists, flooding the victim's mailbox. The objectives of such attacks vary widely—from nuisance and distraction to more serious goals, including obscuring password reset emails within a flood of irrelevant content or priming victims for follow-up impersonation attacks. Based on our analysis of 24 real subscription bombing attack campaigns with 46,970 unwanted emails, we provide insights into operational patterns of these campaigns, examine subscription bombing services offered on the dark web, and discuss mitigation strategies such as automated unsubscription.

CCS Concepts: • **Security and privacy → Network security**; **Denial-of-service attacks**; **Social engineering attacks**;

Additional Key Words and Phrases: email subscription bombing, email flooding

ACM Reference Format:
Mathias Payer and Candid Wüest. 2026. Subscription Bombing: Email Under Attack. *J. ACM* 1, 1 (January 2026), 11 pages. https://doi.org/10.1145/3797487

1 Introduction

Email subscription bombing (also known as subscription flooding or email spam bombing) is an attack technique that overwhelms a target's inbox with a high volume of *benign but unwanted* emails. The emails themselves are benign and are sent by legitimate third parties (such as mailing lists), and therefore typically pass spam filters. The origins of this attack vector trace back to the beginning of the public Internet but have recently re-emerged at scale. Unlike classic email bombing, where attackers send large volumes directly, subscription bombing introduces a layer of indirection. The attacker registers the victim's email address with thousands of third-party online services. These services act as amplifiers that flood the target's inbox with legitimate subscription confirmations and notification messages. Such emails are much harder to block automatically by the receiving email server.

Modern spam filtering relies on multiple signals, including sender reputation, email headers, and message content. For subscription attacks, the sending server generally has a good reputation and is not on a public blocklist; header checks (such as SPF, DMARC, or DKIM) are usually passed, and the subject and body text are different for each email and generally do not match typical phishing phrases that could be detected. While the flood of messages renders the inbox unusable, the true objective is often not just denial-of-service but to hide specific legitimate emails that arrive during the attack, thereby obscuring unauthorized financial transactions, account compromises, or ransomware deployment attempts. Early attacks scanned the internet for sign-up forms and then scripted the sign-up process. In 1996, a stockbroker was bombarded with 25,000 emails, gathering some media attention at the time [1]. Since then, attacks have come in waves and were modernized.

Authors' Contact Information: Mathias Payer, mathias.payer@nebelwelt.net, EPFL, Lausanne, Switzerland; Candid Wüest, candid.wueest@xorlab.com, xorlab, Zürich, Switzerland.

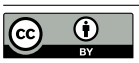

ACM 1557-735X/2026/1-ART
https://doi.org/10.1145/3797487

In 2016, hundreds of email addresses of the US government were targeted, and some were subscribed to over 10,000 newsletters [6]. As organizations increasingly rely on digital communication channels, understanding the mechanics of these orchestrated distractions and their role in broader attack chains has become essential for maintaining operational continuity and financial security.

Beyond newsletters, attackers exploit other services capable of generating emails to a chosen target. This includes password reset requests, account registration confirmations, customer support forms, promotional emails, and social media notifications. These third-party services act as unwitting relays, with attackers hijacking their good reputation and bandwidth to execute the subscription bombing attack. We have monitored email trends since 2015 and have observed a recent increase in such attacks. In this study, we analyze 24 concrete subscription bombing attack campaigns using a large dataset from an email security provider, encompassing 46,970 involved unwanted emails. We define key metrics, outline attack timelines, and provide insight into the operational patterns of these campaigns. Furthermore, we examine subscription bombing services offered on the dark web and categorize their capabilities. Mitigating these attacks is challenging due to the large set of email senders, but automated unsubscription may reduce the impact. Given the simplicity of carrying out these asymmetric subscription bombing attacks, we expect the volume to grow in the future.

2 Subscription bombing

Attackers typically abuse services that are well-known, reputable, and easy to automate using sign-up scripts. Newsletters are the most common targets, followed by account registration confirmation messages. However, attackers have also been seen posting product listings, job postings, support inquiries, hotel bookings, or even apartment listings on online platforms, generating genuine inquiries from human users. The third-party services involved are not compromised or necessarily misconfigured, and no vulnerability in the email server is exploited. The attackers are simply abusing benign (but sometimes overly permissive) features.

The underlying security issue for all services is that, essentially, an attacker can register the victim without their permission who will then receive unwanted messages in their inbox. Key to eradicating this attack vector is that each service will have to improve their sign-up process with a two-step confirmation procedure. There are three key challenges here: (i) a lot of distinct services will have to be updated, (ii) a service provider may not be interested in improving their security if they consider the number of registered users as a success metric, and (iii) even with proper validation the first verification email can still be used to flood victims. As a result, email service providers must mitigate these floods of unwanted emails. Existing academic research on subscription bombing [13] so far targets bypassing filters [9, 14] or detecting unwanted emails through ML-based mechanisms [8].

Some dark web services claim to have access to over 100,000 mailing lists and services. Advances in AI agents may further simplify the process of adding new services to the database of newsletters that can be abused, complete with all the necessary details about which HTTP POST or API request needs to be performed to subscribe new victims. Discussions with supporters of an underground subscription bombing service confirmed that they frequently update their lists with new services.

Subscription bombing is offered as a service but may be run by individual attackers themselves. All an attacker needs is a large list of newsletter providers and an automation script to subscribe the victim's email. An attack consists of a list of victims email addresses that will receive the unwanted emails and a list of sender services which are sending the emails to the victims. A single attack wave consists of the attacker adding the victims address to a preselected number of services within a given timeframe. There can be multiple attack waves in a larger campaign against the same victim. The process of creating a subscription bombing service includes three broad steps:

(1) Build a database of services that send emails by, e.g., crawling the web for newsletters.
(2) Create a script that automatically subscribes new emails to a subset of the database.
(3) Periodically verify the list and replace faulty entries with new ones.

Underground forum discussions indicate that some groups simply record the POST requests required to sign up for newsletters, and then replay it with the new target address. This is not always trivial, as many websites use CAPTCHAs or Cross Site Request Forgery (CSRF) tokens to prevent single-call subscriptions. In the services abused in our attack case study, only 4% used a simple HTTP GET request to subscribe an email address. Around 70% used some form of nonce or CSRF token, requiring an initial HTTP GET request to obtain the token, followed by an HTTP POST request to complete the subscription. By using browser developer tools, an attacker can fill out the required form once, log all requests, and then export the necessary calls directly as a CURL command. This semi-automated process allows attackers to manually build a database. It is also plausible that some attackers targeted WordPress newsletter plugins and similar tools, as these allow them to analyze the subscription process once and replicate it across multiple domains. To test this hypothesis, we analyzed our dataset to identify the Content Management System (CMS) used by the websites, utilizing the service WhatCMS.org for identification. While WordPress was the most commonly used CMS among the sending domains (with a total of 16%, followed by Presta Shop with 6% and Laravel with 5%), we could not identify a single CMS plugin or online shop environment that would indicate a massive monoculture.

With the rise of agentic AI and browser automation tools, attackers can easily commandeer a web browser and have the agent automatically fill in the target email address in subscription forms. Services such as airtop.ai, fellou.ai, or axiom.ai can automate browser interactions and even assist with scaling and CAPTCHA solving. We successfully tested such a setup for newsletter subscriptions. However, to achieve performance and scalability, a paid service would be required. We have not seen any indication that attackers are already using these paid automation services for their attacks. This may change in the future with the evolution of AI tools. Since each newsletter service is hosted by a different web entity, there is no central mechanism to detect or block multiple requests from the same IP address. Nevertheless, most attack services we analyzed claim to utilize proxies to rotate IP addresses and implement rate limitations to avoid getting blocked.

3 Large-scale dataset

Our dataset comprises 46,970 unwanted emails collected across 24 subscription bombing attack waves. The dataset includes email header, To and From address, and subject line. After receiving approval from the involved parties, we gathered the data in the period from October 9, 2024 until December 4, 2024 from midsize and large European customers. Note that collecting emails to form such a study is challenging as the email service provider, the involved company, and individuals have to agree to the data collection. The dataset was collected by the email security solution provider xorlab AG and contains timestamp, sender email address, subject line, anonymized target email address, and extracted tags from email header.

The largest attack wave delivered 4,847 emails in under two hours or 3,387 emails per hour, while the smallest wave delivered 81 emails in one hour. The average number of emails per attack was 1,957 emails or 1,516 emails per hour. The average number of emails received by these accounts on normal days was less than ten per hour. The attackers targeted various industries, including retail, technology, and manufacturing, with no clear pattern. The victims were typically individuals in management positions—director and above—with roles involving external visibility. Their email addresses were publicly accessible on open websites or easy to guess. Victims experienced varying

levels of disruption: 71% of them were targeted once, 18% were hit twice, and 12% faced three separate attacks.

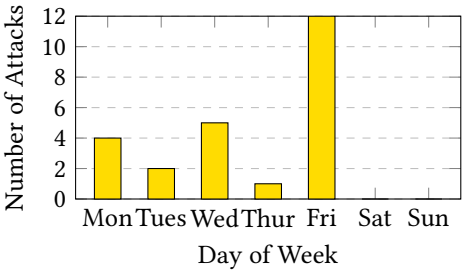

Fig. 1. Subscription bombing attack distribution per week day.

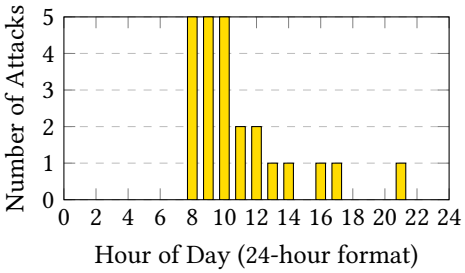

Fig. 2. Subscription bomb attack distribution by hour.

Half of the attacks occurred on Friday (Figure 1), marking the last day of the workweek for these targets. Friday is a common day for cyberattacks, as people may be inclined to leave early, allowing the attacks to go undetected over the weekend. Furthermore IT security teams might operate with limited staff numbers on the weekend. The attacks most frequently began between 8am and 10am (Figure 2), likely to overwhelm victims at the start of the day when they are most likely to check their inbox. 80% of the attacks began during the morning hours of the targets. No attacks were recorded from 10pm to 7am. This suggests that the attackers researched the working hours of their targets for targeted attacks. This timing also aligns with the potential objectives of the flooding attacks, as attackers often need to interact with either the victim or a target service—actions that are not feasible during nighttime hours.

4 Case study: Company Omega

The following analysis is focused on the data from two individual targets of subscription bombing within the same large international organization. Both persons work in upper management and were targeted with one day difference.

- Victim Alpha received 999 emails between Dec 3, 2024 @ 10:19:45 and Dec 3, 2024 @ 11:24:14
- Victim Beta received 719 emails between Dec 4, 2024 @ 12:50:29 and Dec 4, 2024 @ 13:55:08

There was no overlap between the attack waves in our case study, meaning that each victim received emails from different domains. However, we observed minor overlaps across other attacks, suggesting that the pool of available senders is large, with only marginal overlap between randomly selected subsets. The top-level domains (TLDs) of the sending domains vary, with a slight bias towards European countries. This was also reflected in the diverse languages of the newsletter and registration messages. As these emails are automatically generated, 95% of sending email addresses were non-individual accounts (non-human identities). With 26.5% the info@ was the most frequently observed sender account, followed by noreply@ with 7%. Table 1 shows details.

Victim Alpha received 77 account reset emails, while Victim Beta received 49. All 126 services involved were neither financial nor high-profile social media platforms. Instead, they were primarily related to lifestyle and food websites. Therefore, it is unlikely that the attackers attempted to take over any of these accounts for financial gain. Neither of the victims reported any issues with their accounts following the attack. Interestingly, many of these accounts had been freshly created just for the attack. The attacker first created the accounts and then requested a password reset for the newly generated account. The user names of the newly created accounts used random characters like

TLD	Percentage
com	52.79%
de	6.87%
pl	3.14%
org	3.14%
ru	2.68%
net	2.21%
it	2.10%
fr	1.63%
nl	1.57%
es	1.46%

(a) Top10 TLD of sender

Email Prefix	Percentage
info	26.54%
noreply	7.04%
wordpress	6.87%
support	5.59%
no-reply	4.54%
contact	4.07%
admin	2.79%
hello	2.62%
newsletter	1.80%
sales	1.22%

(b) Top 10 account prefixes

Keyword	Percentage
confirm	11.3%
welcome	10.0%
email	8.0%
account	7.3%
password	7.0%
newsletter	6.9%
subscription	6.3%
reset	5.2%
verify	4.1%
address	3.7%

(c) Top10 Email subject keywords

Table 1. Distributions of sender TLDs, account prefixes, and Email subject keywords.

DvqHvlAmxF. Another subscription bombing service improved this aspect and generated accounts with random plausible names like Daniel Williams, Michael Clark, Yoe Miller, or Hannah Fischer. Since the majority of the received emails were newsletters or account registrations, the subject lines commonly contained expected keywords like confirm, welcome, or email. Since the emails received are sent from genuine accounts on legitimate servers, the classic email header verification records like SPF (91.7% pass, 0.3% fail), DMARC (45.2% pass, 3.8% fail), and DKIM (74.3% pass, 12.3% fail) largely passed. Hence, they are not useful for filtering out these attacks. The failure rate is mainly due to badly configured email servers, which is common. Looking at the Spam Confidence Level (SCL), where present, helps to assess if the email was classified as spam and would have been filtered out or would end up in the user's inbox. SCL is a value from 1 to 9 assigned to a message that indicates the likelihood that the message is spam. A value of 9 indicates an extremely high likelihood that the message is spam. A score of 1 is assigned to messages that are deemed not spam. In our dataset, over 70% of the emails received a score of 1 and would end up in the user's inbox if no additional email filtering is used. Around 28% were classified with score 5, which in a default Microsoft 365 configuration would be moved to the Junk folder.

Only 34.4% of the emails did contain an unsubscribe link in the header and only 16.2% contained a one-click unsubscribe link following the header standard for *list-unsubscribe=one-click*. This low number is mostly because we observed many account registration emails and password reset notifications, for which in general there is no unsubscribe link being sent. Furthermore many provider follow the best practice of double opt-in where the user has to confirm their email, before further messages are sent. If we take these into account then we see that we are closer to a 52% rate of emails that require an unsubscribe, of which 36% had an unsubscribe link. The most prominent type of abused website are online shopping websites (25% of our dataset), financial (13% of our dataset), and technology/education (10% of our dataset). These types of sites commonly offer newsletter and account registrations, making them ideal targets. 6.4% of the emails contained some information about where to complain, such as the X-complaints-to header field. Whereas 29.4% had an abuse header like X-report-abuse. Apart from removing the recipient from the mailing list if the user complains, the service provider cannot do much to prevent this.

We conclude that for company Omega, the attackers knew the working hours of victim Alpha and Beta. They used a random subset of their mailing lists to flood them with thousands of emails per hour. Over 71% of the emails were not classified as spam by Microsoft 365 and would have

been delivered to the user's inbox, if no additional filtering solutions would have been in place. Classic header checks (SPF, DMARC, and DKIM) as well as content and link analysis would not have helped prevent this attack. No malicious URLs or malware were involved at any stage. We suspect that the attackers had planned to follow-up with a fake support call to install ransomware, but did not succeed.

5 Goals of subscription bombing

There are multiple potential reasons for attackers to perform subscription bombing attacks against a given target. The most common reasons are: harassment or revenge; distraction from other attacks; pretext for follow-up attacks; or disruption of communication. Even if victims do not immediately recognize the attack as targeted, they quickly notice its impact through the rapidly growing inbox. The primary impact on the victim is the time required to clean up their inbox. Deleting thousands of unwanted emails is time-consuming, and bulk deletion risks removing legitimate messages. Since some of the spam consists of newsletters or subscription-based content, deletion alone is insufficient, more emails will continue to arrive unless the user unsubscribes. Some mailing lists generate high volumes of messages, potentially sending hundreds per day. The actual time spent depends on the email server in use, deployed security solutions, and the responsiveness of the IT team managing the incident. For illustration, assuming 1,000 spam emails, with 80% deleted in one second and 20% requiring 30 seconds to unsubscribe, cleanup would take approximately 113 minutes. Even at an average labor cost of $20/hour, the resulting loss is about $38—compared to the attacker's estimated $1 cost to send the emails.

5.1 Harassment

Individuals with personal grievances may use subscription bombing as a form of digital harassment or retaliation against someone they feel has wronged them. For example, disgruntled employees might target their employer after being denied a desired salary increase. It could also occur in personal relationships, such as an ex-partner. Public figures, such as politicians or CEOs, may be targeted following unpopular decisions. For example journalist Brian Krebs was targeted after some of his discovery articles in 2012 discussed specific cyber criminals [7]. Another example is the US DOGE mail account, which apparently got flooded with spam in early 2025, after requesting status updates from government workers [5]. Such attacks are highly disruptive and frustrating for victims. Victims are often disappointed that built-in spam filters fail to protect them from such attacks [12]. Typically, the goal of harassment is not financially motivated, and the attacker does not directly witness the impact of the attack.

5.2 Distraction

Subscription bombing can serve as a smokescreen to distract from another attack in progress, such as a data breach or account takeover attempt. This is particularly effective when numerous password reset emails are triggered for various legitimate services—among the flood of notifications, a malicious password reset request or suspicious login attempt might go unnoticed, giving the attacker a window of opportunity to exploit it. For example during a noisy credential stuffing attack such a distraction can increase the chances of remaining unnoticed for a longer time period.

Another way attackers profit through distraction is by making fraudulent purchases while burying e-commerce confirmation emails within the bulk of spam. In one case in 2024, a user discovered a purchase confirmation for a new Apple iPhone in the pile of unwanted emails. The attackers used the distraction to carry out a card-not-present fraudulent transaction [15]. They had somehow obtained access to the victim's email and credit card details, possibly from a previous Magecart attack or data breach. Other users on Reddit report digital components being ordered

from Walmart and flights being booked with their email address during a subscription bombing attack. The primary goal of subscription bombing in such cases is to conceal the real attack for as long as possible.

5.3 Follow-up Attack

Instead of using subscription bombing solely as a distraction, attackers can also exploit it as a stepping stone for more targeted deception. There have been multiple reports of threat actors following up a subscription bombing campaign by impersonating the victim's IT helpdesk. This impersonation typically takes place via alternative communication channels such as Microsoft Teams calls or WhatsApp messages.

Under the disguise of helping resolve the spam overload, attackers persuade the victim to initiate a remote access session, thereby unknowingly granting control to cyber criminals. Commonly used remote monitoring and management tools, such as AnyDesk, are often employed to facilitate this access. In 2024, Rapid7 reported several incidents in which the Black Basta ransomware group used this tactic to gain remote access and deploy additional malware [4, 10]. Leaked internal chat logs from the Black Basta ransomware group from February 2025, revealed this approach as a part of their playbook. One message read: "*I will flood their inboxes with spam, and you will call them pretending to be an IT admin, saying they need to install a spam filter. She installs AnyDesk, and we get in to install our software.*" [3]. The primary objective is to create a scenario in which the victim is more likely to accept external assistance, ultimately leading to system compromise.

5.4 Overloading

Attackers may target organizations to overload their email servers, causing technical problems or outages within the network similar to a DDoS attack against the web server. Depending on the configuration and setting of the target email server and inbox, there may be a threshold beyond which even legitimate emails are rejected by the target. This allows the attacker to temporarily block important emails from even reaching the inbox, until the resources have been freed and a retransmission is attempted [2]. In 2025 the hacktivist collective Anonymous launched attacks to disrupt the online activities of ISIS [11]. Among DDoS and defacement attacks this included email flooding. This could disrupt the normal work of the employees during critical time intervals and is frequently combined with an extortion attempt. The attackers ask for money in order to stop the attack. Similar to network DDoS attacks, the victim might give in and pay the blackmail.

6 Subscription bombing as a service

Various cyber criminals offer subscription bombing on the dark web and underground web sites, often with a free trial. The criminals typically boast that they can flood the target inbox with 1,000-10,000 emails per hour with a spam detection rate of less than 0.1%. One service claims that they can generate 6,000 emails per hour with over 100,000 newsletters and forums in their database. Generating such a service is not very difficult on a technical level. Analyzing the leaked chat messages from the Black Basta ransomware group [3] shows that they created their own subscription bombing service. Subscription bombing services are advertised using names like "emailbomb.cc", "ebomber.cc", "floodbox.org", "diddy.pw", "subscribethemall.com", or "mailbait.info". Note that attacking inboxes without consent is illegal in most jurisdictions.

As a case study, we explore one representative subscription bombing service, accessible via a publicly available website on the regular internet. The platform primarily offers subscription bombing, though SMS/text message bombing is also available. Attacks against a specific target can be scheduled for a duration ranging from a minimum of 1 hour to a maximum of 31 days. Emails can be sent using a "drip feed" model, distributing the load over time, with adjustable rates from 75

to 3,000 emails per hour. Pricing is straightforward: 1,000 emails can be sent for as little as $1 using the standard list. A "high quality" list, marketed as having a higher chance of bypassing spam filters, is available for $10 per 1,000 emails. The service dashboard provides an overview of all scheduled tasks and their completion status. Communication is handled through a Telegram channel, and support is also available via an integrated AI chatbot (tawk.to). For advanced users, the provider also offers a REST API to automate attacks further. Payments are handled exclusively through cryptocurrency via a gateway provider (pay.crypto-gateway.xyz), which currently supports BTC, LTC, and XMR.

In our own test using their free trial with the standard list, 21% (14/68) of the emails arrived in the inbox after 13 minutes. A second trial had 24% (929/3,930) of emails reaching the inbox. Note that attacking inboxes without consent is illegal in most jurisdictions. We only tested mailboxes directly under our control, on our servers.

By extrapolating from the order numbers in the dashboard, it appears that at least 300,000 tasks have been executed to date, assuming the count started at zero. During a test window of 767 hours, the task counter increased by 17,323. That represents 22.6 campaigns per hour. The activity level indicates ongoing interest for such services.

Table 2 shows prices for subscription bombing campaigns, taken from service offers across different underground websites. One of the scripts with unlimited access was available to buy for $500. The prices are very similar. The major variation is the quality of the list, which results in a higher inbox success rate, at a higher price tag.

Service	Service A	Service B	Service C	Service D
Free trial	10 min	25 emails	25 emails	100 emails
5,000 emails	1 hour $9	$5	$50	$5
10,000 emails	1 week $65	$10	$100	$10

Table 2. Prices to send N emails in a subscription bombing campaign.

7 Mitigation

As subscription bombing is increasingly used to target individuals and organizations, defenders must apply layered countermeasures. To minimize impact, organizations should:

(1) set thresholds or rate limits for the number of emails an inbox can receive from previously unseen senders,
(2) educate users about subscription bombing attacks,
(3) group similar newsletter or account reset emails during active attacks, and
(4) automatically unsubscribe or bounce unwanted newsletter messages where supported.

Additionally, the receiving server may flag incoming subscription bombing attacks and alert the security team.

While securing the account itself through two-factor authentication (2FA) or zero-trust access (ZTA) is recommended, these measures do not mitigate subscription bombing. All the attacker needs is the victim's email address, which could come from a previous data breach, or in most cases has been identified from social media or because of a predictable naming schema such as firstname.lastname@company.tld. When interacting with diverse services on the Internet, the email address naturally "leaks". Some users create individual email addresses when registering with online services by appending the plus sign (e.g., *yourname+newsletter@email.tld*). While this allows the user to pinpoint where the email address leaked, it does not protect against the attack.

The same holds true for security awareness training through phishing simulations. While those are useful for security in general, they cannot mitigate subscription bombing attacks.

A cleanup process based exclusively on keyword blocking is tricky, as this would also remove legitimate and wanted newsletters and emails. Moving all suspicious emails to a different folder and have the user manually inspect them is not a viable option, as this would still require user time to sieve through thousands of emails.

7.1 Newsletter Providers

As newsletter and service providers are being abused as amplifiers in these attacks, they are hopefully interested in hardening their services. Providers of newsletters can limit abuse of their services with the following measures. Newsletter subscription should require an email confirmation (double opt-in) and single-click unsubscribe. Having the user click on a link they receive to confirm subscription ensures that the newsletter was delivered and received correctly. To prohibit automated registrations, the sign-up page should require a CAPTCHA, even if their benefits are decreasing with new AI tools. Another angle is to require users to register an account before being able to subscribe to slow down the registration process. Even though several of these recommendations add friction to the registration process, service providers should be encouraged to use them. Otherwise, service providers risk being added to blocklists, e.g., when victims report the newsletters from attacks.

7.2 Filtering and Triaging

The email server could build a list of known-good addresses at the server level, domain level, and account level. These lists can then be used to tag known-good emails during an ongoing wave to help users mitigate negative effects and allowing them to process the flood, if not automatically, asynchronously. Curating a trusted list by adding domains based on interaction is impractical, as users typically do not reply to newsletters, making it a one-way communication pattern. This leaves only an age score for newsletter domains from which the user has previously received emails.

7.3 Unsubscribe

Unsubscribing after a subscription bombing attack is complex. Deleting all emails within a given time frame may also remove legitimate, important messages. Furthermore, deleting alone does not solve all issues. Since many of these emails originate from newsletters, unsubscribing is necessary to prevent continued flooding. But this action is very time consuming if done manually. Some suggestions from the Internet, such as changing the email address, are impractical in most cases. Some newsletter services make unsubscribing unnecessarily complicated such as requiring CAPTCHA solving, account creation, or asking about communication preferences and feedback. Also, some providers simply ignore changes made by the user and continue to send emails.

These obstacles make automated unsubscription difficult and, in some cases, even risky. Unfortunately, this does only help with newsletters and not with other types such as password reset emails which do not have unsubscribe options. As alluded earlier, there are many cases where the services used are non-English, making it harder for humans to find the unsubscribe option in the various languages. Depending on the email server's capabilities the administrator can set up blocks or bounce messages for involved emails. Unfortunately this does not guarantee that the sender will refrain from sending more emails in the future. GenAI can help with unsubscribing from such services by interacting with the page and finding the right option. In our tests with early agentic browsers the success rate of AI-supported unsubscription was less than 50% but recent improvements in AI agents show promise in improving the success rate of agentic unsubscribe tools. Further research in automatic unsubscribing is necessary, especially after detecting an attack.

8 Conclusion

Subscription bombing attacks are technically simple for attackers to conduct yet disproportionately difficult for victims to mitigate. Based on our analysis of 24 subscription bombing campaigns, this emerging attack vector poses a growing operational and security risk to digital communication. The attack itself provides only limited use-cases for direct fraud, but can be used as distraction to hide concurrent attacks. Attackers leverage subscription bombing for various malicious purposes, ranging from creating nuisance to obscuring critical security alerts and facilitating social engineering attacks. Our study provides valuable insights into the operational patterns, timelines, and dark web services associated with these campaigns.

While complete mitigation remains challenging due to the asynchronous distribution between attackers and defenders, automated unsubscription offers a promising avenue for reducing the impact on victims. Legislative options such as requiring unsubscribe links are another option to reduce the impact of subscription bombing. Further research is crucial to develop more robust defense mechanisms against this evolving form of digital attack.

9 Acknowledgments

We thank Max Fischer for data collection and members of xorlab and HexHive for feedback. This work was supported in part by the ERC (grant No. 850868) and the SNSF (grants PCEGP2 186974 and 200021-236559).

References

[1] Seymour Bosworth, Michael Kabay, and Eric Whyne. 2014. *Computer Security Handbook, Section 18.5.2.* Wiley, Hoboken, NJ, USA. https://books.google.ch/books?id=yKQ6AwAAQBAJ

[2] Health Sector Cybersecurity Coordination Center. 2024. HC3 warns about email bombing in health care sector. https://www.hhs.gov/sites/default/files/email-bombing-sector-alert-tlpclear.pdf.

[3] D4RK-R4BB1T. February 24, 2025. Leaked chat logs from Black Basta ransomware group. https://github.com/D4RK-R4BB1T/BlackBasta-Chats/blob/main/blackbasta_chats.json.

[4] Derek B. Johnson. 2025. Ransomware groups pose as fake tech support over teams. https://cyberscoop.com/ransomware-groups-pose-as-fake-tech-support-over-teams/.

[5] Jordan King. February 25, 2025. Elon Musk's DOGE Flooded With Spam After Email Threat to Workers. https://www.newsweek.com/elon-musk-doge-spam-threat-workers-2035702.

[6] Brian Krebs. August 16, 2016. Mail bombing .Gov email addresses.

[7] Brian Krebs. July 18, 2012. Cyberheist Smokescreen: Email, Phone, SMS Floods. https://krebsonsecurity.com/2012/07/cyberheist-smokescreen-email-phone-sms-floods/.

[8] Aurobinda Laha, Md Tahmid Yasar, and Yu Cheng. 2022. SubStop: An analysis on subscription email bombing attack and machine learning based mitigation. *High-Confidence Computing* 2, 4 (2022), 15 pages. doi:10.1016/j.hcc.2022.100086

[9] Jinrui Ma, Lutong Chen, Kaiping Xue, Bo Luo, Xuanbo Huang, Mingrui Ai, Huanjie Zhang, David S.L. Wei, and Yan Zhuang. 2024. FakeBehalf: imperceptible email spoofing attacks against the delegation mechanism in email systems. In *Proceedings of the 33rd USENIX Security Symposium*. USENIX Association, Philadelphia, PA, 1243–1260.

[10] Tyler McGraw Thomas Elkins Evan McCann. 2024. Black Basta Ransomware group uses email bombing as distraction. https://www.rapid7.com/blog/post/2024/05/10/ongoing-social-engineering-campaign-linked-to-black-basta-ransomware-operators/.

[11] Simon Parkin. October 6, 2016. Operation Troll ISIS: inside Anonymous' war to take down Daesh. https://www.wired.com/story/anonymous-war-to-undermine-daesh/.

[12] Brian Robbins. April 19, 2023. Help from Google Support after email bombing. https://support.google.com/mail/thread/211841254/just-got-email-bombed-today-and-gmail-did-not-warn-me-like-it-says-it-s-supposed-to-has-anyone-else?hl=en.

[13] Markus Schneider, Haya Shulman, Adi Sidis, Ravid Sidis, and Michael Waidner. 2020. Diving into Email Bomb Attack. In *50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, Piscataway, NJ, 286–293. doi:10.1109/DSN48063.2020.00045

[14] Kaiwen Shen, Chuhan Wang, Minglei Guo, Xiaofeng Zheng, Chaoyi Lu, Baojun Liu, Yuxuan Zhao, Shuang Hao, Haixin Duan, Qingfeng Pan, and Min Yang. 2021. Weak Links in Authentication Chains: A Large-scale Analysis of Email

Sender Spoofing Attacks. In *Proceedings of the 30th USENIX Security Symposium*. USENIX Association, Philadelphia, PA, 3201–3217.

[15] Katherine Wood. 2024. Email bombing distracting from fraudulent purchases. https://www.signifyd.com/blog/email-bomb-spam-attacks/.