SpamRATS!

Really Annoying Trouble Spots

SpamRATS!
Really Annoying Trouble Spots

Introduction

Mail server operators face the challenge of securing communication for Internet users and are in a unique position where they not only need to properly classify legitimate marketing email, but also thwart the ever increasing threat activity from bad actors and criminals. Email accounts are required to enjoy the many applications and services the Internet has to offer, and the increasing adoption and integration of the Internet to everyday life further increases the value of these accounts. This catches the eye of email marketers and bad actors, overwhelming mail servers with unwanted email and authentication attacks in order to access this value. Effective mitigation of the unwanted email and authentication attacks that are threatening mail servers can be achieved with IP Reputation.

Unwanted email

One of the major problems that email systems face is unwanted email. "Spam" has a very specific definition for legal purposes, and oftentimes this definition would exclude emails the average person would call Spam. So instead, we'll use the term unwanted email, which refers to email that the majority of people would realistically not want (such as unsolicited marketing, phishing, or malware). Unscrupulous email marketers and bad actors are behind the unending sea of unwanted email, and although the lawfulness behind their actions may differ the intended outcome is to make money.

Authentication Attacks

Another major but often not talked about problem that email systems face are authentication attacks. Authentication attacks are all of the different kinds of password guessing methods that commonly plague servers. All sorts of bad actors are behind these attacks, from criminal groups selling access to accounts to state-sponsored hackers conducting reconnaissance for their employers.

IP Reputation

The threats mentioned have something in common: a connection is required for the threat to be possible. In other words, an IP address will need to connect to a mail server, and this is where IP Reputation plays a role. IP Reputation can identify whether a connection is safe, or aid in ensuring that the connection falls within a specific criteria before allowing it.

SpamRATS!
Really Annoying Trouble Spots

Table of Contents

SpamRATS!
Really Annoying Trouble Spots

Unwanted Email

How much Email is Unwanted?

When looking at the statistics on the percentage of total email considered as unwanted, numbers range from 40% to over 80%. This number depends on how unwanted email is calculated. If it is calculated by referring to the total number of emails that are successfully delivered to email accounts (both in the inbox and the junk folder), the percentage of unwanted email may be at the lower end of the spectrum at around 40%.

“On the percentage of total email considered as unwanted, numbers range from 40% to over 80%”

However, if the calculation is done by referring to the total number of attempts to send email to a mail server, the number of attempts to deliver unwanted email can skyrocket to 90%. Of course, some email systems perform better than others when it comes to mitigating unwanted email. What we can conclude is that a significant portion of email is unwanted, but why does that matter?

Problems that Unwanted Email Creates

In order to send an email, the server, machine, or device that is sending the email needs to connect to the mail server that it is trying to send to. If over 80% of these connections are considered unwanted, that means a mail server needs to sift through all that junk in order to identify the legitimate emails. Unwanted connections can take up a lot of a mail server's resources, even to the point where the mail server stops functioning properly. This could result in the mail server taking longer to send and receive emails, or even losing emails altogether.

On the side of the email users, unwanted email can be annoying to potentially devastating. It can result in time wasted by having to manually identify and remove unwanted email, and those minutes add up over time. Phishing emails are cases where the sender pretends to be someone they're not in order to convince the recipient to do something they otherwise would not have done. This can trick people into giving up private information or downloading malware, increasing the chance of stolen accounts, compromised machines, and even monetary losses.

SpamRATS!
Really Annoying Trouble Spots

Businesses have a lot more to lose than individual email users. The FBI's Internet Crime Complaint Center (IC3) reported business email compromise (BEC) causing losses of $1.8 billion USD in 2020[1]. BEC refers to different methods involving email that criminals use in order to steal or extort money from businesses. One example is sending a fake email invoice pretending to be a well known company in hopes of fooling the recipient. A fooled recipient could send money to the criminals bank account or download malware that the criminal sends. Malware could lead to ransomware (a type of malware that locks access to the files of a computer) being installed onto computers, crippling and exposing the business for extortion.

Authentication Attacks

Authentication Attacks are getting more Sophisticated

Brute force attacks are likely the first thing that comes to mind when thinking about authentication attacks. Brute force attacks refer to password guessing attempts that try every possible combination of letters, numbers, and symbols in hopes of gaining access to the account. Today there are much more effective methods being used. Dictionary attacks use predetermined lists of words for guessing passwords. You can bet that the most common passwords[2] will be attempted, they're the most common for a reason!

Credential stuffing attacks are becoming more common. These attacks leverage the information leaked in data breaches, which are cases where businesses have their private data exposed to unwanted eyes, caused by poor security settings or bad actors ex-filtrating data from their systems. These breaches may reveal usernames, passwords, sometimes even both. This information is then used in credential stuffing attacks, and they are extremely effective due to the common habit of reusing passwords across multiple accounts.

Problems that Authentication Attacks Create

Many applications and services on the Internet require an email account to sign up, and in some cases merely having access to the associated email account is enough to access these services. Password reset tools are commonly tied to email accounts, allowing anyone with access to the email account to access the service. Regaining control over their email accounts is a frustrating experience for end users and uses up Support's time and resources.

---

1 <https://www.ic3.gov/Media/PDF/AnnualReport/2020_IC3Report.pdf>
2 <https://en.wikipedia.org/wiki/List_of_the_most_common_passwords>

Businesses again have a lot to lose. Email accounts can contain important and private emails which bad actors can use to their advantage. Proprietary company information could be ex-filtrated, private emails exposed, and important events intercepted. For example, a bad actor can observe that a large payment is about to be made and utilize the compromised account to redirect that payment into a different account. This is another way BEC[3] can occur.

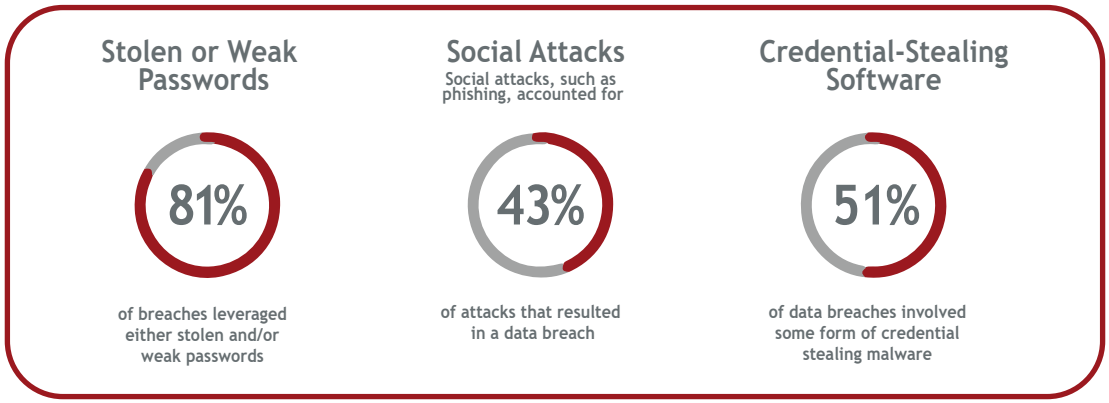

https://www.okta.com/resources/whitepaper/5-identity-attacks-that-exploit-your-broken-authentication/

IP Reputation

IP Reputation is a method that identifies whether an Internet Protocol (IP) address is suitable, or safe, to accept a connection from. For example, if you only want IP addresses located in Canada to connect to your server, you could get the list of IPs containing all the IP addresses located in Canada and use it to check the IPs that connect to your server. If the IP matches with one in the list the connection is accepted. You can do the reverse, and reject connections with IP addresses on that list, essentially blocking Canadian IP addresses from connecting to your server. In other words, IP Reputation lists are lists that contain IP addresses that all fall under a specific set of criteria.

This criteria is usually related to whether or not connections should be accepted. For any server's individual use case, understanding the criteria for connections you want to allow and connections you want to block is important when deciding what kind of IP Reputation list to utilize.

3 For more information on BEC <https://www.rcmp-grc.gc.ca/en/business-email-compromise-bec>

SpamRATS!
Really Annoying Trouble Spots

SpamRATS specializes in creating IP Reputation lists for the purpose of mitigating the threats that mail servers face. These IP Reputation lists are provided in the form of Real-Time Blackhole lists (RBLs).

The criteria used for the SpamRATS reputation lists reflect what we believe are 'Best Practices'[4] for mail server operators.

This whitepaper goes over four IP Reputation lists that SpamRATS offers:

- **RATS-NoPtr**
- **RATS-Dyna**
- **RATS-Spam**
- **RATS-Auth**

4 Further reading on 'Best Practices'
<https://www.m3aawg.org/sites/default/files/document/M3AAWG_Senders_BCP_Ver3-2015-02.pdf>

SpamRATS!
Really Annoying Trouble Spots

RATS-NoPtr

The RATS-NoPtr reputation list contains IPs that do not have a reverse DNS record (rDNS, also referred to as 'hostname' or 'PTR' record). Referring to 'Best Practices', the IP of a properly configured mail server should have a rDNS that reflects the domain of the party responsible for any emails being sent by that server.

The criteria for additions to the RATS-NoPtr reputation list is:

- **At the time of detection, the IP did not have a rDNS configured (shows up as 'NXDOMAIN')**
- **The IP attempted to deliver unwanted email, or to non-existent email accounts.**

Seeing this activity across the same network suggests that the provider for those IPs allow email to be sent from their network, something that can be disabled if they did not want just anybody to have that capability.

Aside from misconfigured mail servers, many of the IPs with no rDNS record are machines infected with malware capable of sending email. Bad actors, people involved in malicious internet activity, control these compromised machines often using them to send unwanted email. Why don't they just update the rDNS record to improve their chances of successfully sending email? In most cases they can't, because it requires the person assigned the IP to contact their provider in order to get the rDNS updated.

Our stance is that if an IP without a rDNS record attempts to deliver email, it should not be accepted. Before sending email, a mail server should be properly configured, and having a rDNS record is an important part of the configuration. This significantly reduces the amount of emails received from compromised machines. Once the rDNS is updated, the mail server operator can easily remove themselves from RATS-NoPtr.

SpamRATS!
Really Annoying Trouble Spots

RATS-Dyna

The RATS-Dyna reputation list contains IPs that have a rDNS matching the pattern seen in dynamic networks. Dynamic networks are networks that assign IPs to machines in a dynamic fashion. For example, a computer on a dynamic network could have one IP address today, and a different one tomorrow. Referring back to 'Best Practices', the IP of a properly configured mail server should have a rDNS that reflects the domain of the party responsible for any emails being sent by that server. IPs that have a rDNS matching a dynamic pattern usually contain the provider's domain rather than the responsible party's domain.

The criteria for additions to the RATS-Dyna reputation list is:

- **At the time of detection, the IP had a rDNS that matches a dynamic pattern (e.g. dynamic.11.11.11.11.myprovider.com)**
- **The IP attempted to deliver unwanted email, or to non-existent email accounts**

Similar to the RATS-NoPtr list, seeing this activity across the same network suggests that the provider allows dynamic IPs to send email, despite having the ability to disable this. A properly configured mail server should always be using a static IP. If a dynamic IP is used for a mail server, communication with that mail server becomes unreliable. For example, when replying to an email sent from a mail server with a dynamic IP, the mail server's IP could have already changed and the reply would fail to send to that server.

Mail operators know that dynamic IPs aren't suitable for mail servers, which means that the majority of email connections from these IPs are by compromised machines. It isn't enough to manually block these IPs either, because the same compromised machine could connect again from a different IP.

Our stance is that emails from dynamic IPs should be filtered to the junk folder. Networks having rDNS records configured would suggest that the provider is diligent and active in clearly defining the usage of their networks. In contrast to RATS-NoPtr, there's a slim chance that a legitimate mail server tried to send email before updates to their rDNS have fully propagated, so filtering to the spam folder is a good compromise. Once the rDNS is updated, the mail server operator can easily remove themselves from RATS-Dyna.

SpamRATS!
Really Annoying Trouble Spots

RATS-Spam

The RATS-Spam reputation list contains IPs detected as sending high volumes of unwanted email or sending to many email addresses which do not exist within a short period of time. These IPs have custom rDNS records, suggesting that they are purposed for sending email and do not belong in the previously discussed reputation lists. A lot of this activity comes from networks that have a history of allowing this activity and not appearing to take any steps towards proactively mitigating it. We've also taken extensive precautions to avoid adding IPs from services that are considered "too big to block" such as Freemail services, or entities that provide important communications.

The criteria for additions to the RATS-Spam reputation list is:

- **At the time of detection, the IP's rDNS does not qualify for RATS-NoPtr or RATS-Dyna**
- **The IP attempted to deliver high volumes of unwanted email or non-existent email accounts across multiple unique ISPs (Internet Service Providers) within a short period of time**
- **The IP may belong in a network with a history of sending unwanted email**
- **The IP is detected as sending malicious email such as phishing and malware.**

There are many valid reasons a legitimate server could be listed. They could have an issue with compromised accounts, or the email server itself could be compromised and in the control of bad actors. It is common for any mail server to leak a little bit of unwanted email, but high volumes over a period of time targeting multiple unique ISPs is unacceptable. The operator of a mail server that has been listed on RATS-Spam should perform an audit of their mail server, because there is a good reason that it was added in the first place.

Many Internet and Hosting providers either allow unwanted email to be sent from their networks, or are unable to mitigate the abuse coming from their network. SpamRATS' technology automatically adds IPs exhibiting abusive patterns the moment it is detected on our data collection grid of specialized mail servers.

Our stance is that being listed on RATS-Spam isn't about being punished for sending unwanted email; it is about being made aware that there is an issue that needs to be addressed with your mail server. That is why we've made it extremely simple for the public to remove their IPs from RATS-Spam. Responsible mail operators should fix the issue before removing their IP. Irresponsible or negligent mail operators that continue to allow their mail systems to be abused will get relisted.

SpamRATS!
Really Annoying Trouble Spots

RATS-Auth

The RATS-Auth reputation list is distinct from the previous SpamRATS reputation lists, as it contains IPs that have been detected performing authentication attacks. It should be used to detect suspicious login attempts, rather than unwanted email. This list is comprised of static IPs from networks that either allow, or are unable to mitigate, the abuse of their networks to perform password guessing attacks.

The criteria for additions to the RATS-Auth reputation list is:

- **The IP has been detected as performing authentication attacks**
- **The IP is from a network that is known to be a source of authentication attacks**
- **The IP is static**

Proper precautions are made to avoid adding dynamic IPs or Carrier Grade NAT (CGNAT) networks. Adding dynamic IPs could affect legitimate end users from logging in, as they would be punished for the activity that the previous owner of the IP was responsible for. Likewise for CGNAT networks there could be many end users behind the same IP.

Mail operators may find it difficult to effectively mitigate the authentication attacks brute forcing their mail servers. Not only is there a lot of unrelated noise in the mail logs to sift through, but many considerations need to be made before the right IP can be identified as safe to block. RATS-Auth can be used to ease this burden, as our criteria for additions makes it extremely safe to use.

Our stance with RATS-Auth is that it is the responsibility of the network providers to identify and mitigate the abuse coming from their networks. It doesn't seem to be a resource issue, as we see authentication attacks coming from the networks of some of the biggest companies in the world. These attacks come in high volume, from new IPs daily, with very consistent patterns. It feels unreasonable to put the burden of mitigating this threat on mail operators, considering where the threat is coming from.

SpamRATS!
Really Annoying Trouble Spots

Concluding Remarks

IP Reputation is a powerful tool that can mitigate the threats that mail servers commonly face such as unwanted email and authentication attacks. A good understanding of the nature of these threats for your particular circumstances allows you to minimize false positives when using solutions such as IP Reputation lists. However, it isn't meant to be a complete solution. IP Reputation provides the ability to quickly mitigate common threats within a specific criteria, freeing up resources to tackle more difficult issues that require more complex solutions.