Detecting Gray in Black and White

Christian Rossow, Thomas Czerwinski, Christian J. Dietrich, Norbert Pohlmann
Institute for Internet security
University of Applied Sciences Gelsenkirchen, Germany
{rossow|czerwinski|dietrich|pohlmann}@internet-sicherheit.de

ABSTRACT

DNS based black- and whitelists are heavily used in the fight against spam. However, in certain cases their use can cause conflicts, such as *false positives*. In our work, we show a method to identify IP addresses that are listed in both black- and whitelists. We term this set of addresses as *gray IP area*. We then develop a method to classify these senders as either spammers or legitimate mailers. This method is applied in an experiment using well-known black- and whitelists. The results show difficulties in such an automated classification because of legitimate mail servers relaying spam mails, mail senders behind dynamic IP addresses and misconfigured MTAs. We conclude that there is no automated mechanism to perform a reasonable classificiation without manual expert knowledge of all involved mail senders.

1. INTRODUCTION

Since the early ages of email, DNS based blacklists have been widely used as anti-spam measure. They allow to filter or deny SMTP traffic from senders using IP addresses with a bad reputation. In particular high-volume mail receivers tend to use blacklisting, since it is a very efficient filtering mechanism that does not require to inspect the mail content at all. In addition, professional services conveniently operate high-quality DNS based blacklists and supply these to mail receivers. Thus, hitherto, blacklisting is *the* (or at least one of the main) mechanisms used by email operators to filter unsolicited messages.

However, blacklists have always been prone to false positives. In the context of blacklists used as anti-spam tool, false positives are entries that wrongly lead to blocking legitimate mail submissions. In order to soften that risk, blacklists are usually combined with whitelists. Such whitelists include IP addresses or mail servers that are considered to send legitimate mail - even accepting a certain spam ratio. As a consequence, these *cross-listed* mail servers may always initiate SMTP connections to submit mails and a spam check is done at later stages, if at all. Experiences of mail experts showed that the mechanism works well in most cases. But if the interaction of black- and whitelists fails, legitimate mail senders can be blocked from mail transmission until further manual interference.

Although the coarse granularity of blacklists has been known to anti-spam experts for years, to the best of our knowledge, we are the first to explore cross-listed IP addresses systematically. We see a need for a more detailed discussion due to recent and current spam trends that partially undermine IP reputation based anti-spam mechanisms. First, in the early beginnings of spam the unsolicited mails were sent by open relays with fixed IP addresses. Nowadays the majority of spam is sent by zombie computers that are part of a botnet with dynamically assigned IP addresses [5, 7]. This makes it difficult for blacklists to have a complete set of spamming IP addresses [14, 11, 6]. Next, current malware tends to steal data from home users that also include credentials of legitimate SMTP servers. The stolen credentials are then used by spammers to relay unsolicited mails through SMTP smarthosts that otherwise send mostly legitimate mail. Third, malware can even sign up for new accounts by evading CAPTCHA mechanisms of free mailers [10]. These accounts can also be misused for sending spam from an else well reputed mail source. As a last option of this still incomplete listing, a conflict occurs if users instruct a primary mail account to automatically forward all received mail to a secondary mail account located on a different mail server. In this situation, not only legitimate mail but also spam is typically forwarded to the secondary account and contaminates the primary mail server's reputation.

In this work, we perform a detailed analysis of conflicts that may arise when using white- and blacklists to combat spam. First, we develop a metric that we call *intersection matrix*, which shows the relation between different black- and whitelists. It reveals many intersections among DNS based lists. We particularly focus on those IP addresses that are listed in both black- and whitelist and term these set of addresses *gray IP area*. We find that, globally, every 6300th SMTP connection is established from such a cross-listed IP address. Given this high impact of potential false positives, we introduce a method to further inspect the gray IP area. We conclude that there is no reasonable automated possibility to classify these senders as either good or bad. As a consequence, we substantiate a fact which was postulated by many email operators before: Blacklisting is an efficient anti-spam mechanism, but is becoming more and more prone to false positives.

2. CROSS-LISTINGS

In this section, we will discuss intersections among multiple black- and whitelists. We will show that lists often have some IP addresses in common. Our discussion will then focus on those IP addresses that are part of both a black- and a whitelist, the so called gray IP area.

MIT Spam Conference 2010 March 25-26, 2010, Cambridge, MA USA

2.1 DNS based reputation lists

Reputation based lists are commonly applied by anti-spam experts to filter SMTP connections from senders that send spam. Whereas blacklists can be retrieved from professional blacklist operators (such as e.g. Spamhaus.org [13, 8]), whitelists are typically managed by email operators themselves. In this work, we investigate the interplay of four well-known high-quality blacklists with a whitelist operated by a big internationally operating anti-spam service company. We monitored these lists over a period of 20 days from 2009/08/07 until 2009/08/26 and performed daily analyses. Table 1 summarizes the average size of the monitored lists.

list	*entries*	*IP range*
NiX Spam	338,551	338,551
Spamhaus SBL	5,241	1,968,208
Spamhaus PBL	1,068,533	552,177,454
Spamhaus XBL	7,304,867	7,304,867
Reference whitelist	11,936	11,936

Table 1: Sizes of lists measured in listings and covered IPv4 address range.

The two Spamhaus lists SBL and PBL share a particular property that do not apply to the other monitored lists. Whereas most lists include single IP addresses (/32s) only, Spamhaus' SBL and PBL list entire IP network ranges that should not send legitimate mail. As a first measurement, we cross these lists and create a matrix that indicates which IP addresses of list A are also part of list B. Figure 1 illustrates an example matrix from 2009/08/25.

reference / comparison	NiX Spam	sbl.spamhaus.org	pbl.spamhaus.org	xbl.spamhaus.org	reference whitelist
NiX Spam	-	0.1	81.7	89.3	0.0
sbl.spamhaus.org	0.0	-	4.4	0.1	0.0
pbl.spamhaus.org	0.1	0.0	-	1.2	0.0
xbl.spamhaus.org	4.4	0.0	93.5	-	0.0
reference whitelist	0.6	0.1	4.7	0.7	-

Figure 1: Intersection matrix of selected lists

The matrix shows which ratio of IP addresses of the list in a row is covered by the list in the column. Obviously, the matrix shows intersections between different blacklists. Spamming mail servers may be listed in several blacklists. In particular if the ratio of legitimate mail originating from an IP address is very low and its spam ratio is very high, it is often listed in multiple blacklists. For example, the NiX Spam blacklist is covered by the Spamhaus PBL blacklist to 81.7%.

Similarly, whitelists often share the same IP adresses of legitimate senders. Interestingly, there are several cases where an IP address is listed in both, black- and whitelists. We call such an IP address *cross-listed* and define the set of cross-listed IP addresses as *gray IP area*. Thus, the gray IP area consists of IP addresses that are special cases of listings. At first sight, any cross-listing indicates either an erroneous listing on a blacklist or an erroneous listing on a whitelist. But at a closer look, there are many reasons why IP addresses become listed on both types of lists simultaneously:

- The input procedure to black- and whitelists, i.e. if an IP addresses shall be listed, may conflict. If, for example, a blacklist relies on manual user feedback, users may erroneously submit legitimate newsletters as spam, whereas a whitelist classifies this bulk sender as legitimate.

- Freemail providers often have issues with fraud users that misuse the free service. For example, it has been observed that malware is able to evade CAPTCHA mechanisms of freemail providers [10]. The malware can then register new user accounts and send spam via these accounts. Although in total most mail sent by such a freemail provider is legitimate, some users abuse the free service to submit spam.

- Furthermore, malware is capable of stealing credentials of legitimate SMTP servers that are stored at home user computers [15, 1, 2]. The stolen credentials are then used by spammers to relay unsolicited mails through SMTP servers that otherwise send legitimate mail only.

- Conflicts may occur if users can configure a primary mail account to forward all received mails to a secondary mail account at a different server. In particular if received spam is not filtered by the primary mail account, but forwarded to the secondary account, the first mail server appears as a spam source and may be listed on a blacklist.

- The reputation of IP addresses is not fixed and may change, if for example the responsible entity behind the IP address change [12]. In this case, old listings of either black- or whitelast can become outdated.

Therefore, cross-listings are a reasonable implication of mail senders that send both legitimate and unsolicited messages. Thus, a deny-all or allow-all policy for these senders is prone to many wrong decisions [9]. An individual check of each SMTP connection or message, respectively, prevents from lumping together all mails from a specific sender. In this section, we will give an estimate of the global impact of cross-listings.

2.2 Gray IP area

We defined the gray IP area as the set of all cross-listings, i.e. IP addresses that are included in both our whitelist and at least one blacklist. During the test period, the gray IP area size of each daily snapshot varied between 654 and 913 IP addresses. On average, 750 distinct IP addresses were cross-listed. Thus, a fraction of 6.3% of our reference whitelist was listed in at least one of the blacklists. Table 2 shows how much every single blacklist contributes to this percentage. Note that the sum of all intersections exceeds 6.3% due to the filter of distinct IP addresses.

A huge fraction of the intersection was caused by listings of the Spamhaus PBL. Whereas most of the blacklists list single IP addresses only, the PBL is a special case. It

blacklist	*WL by BL*	*BL by WL*
NiX Spam	0.42%	0.0142%
Spamhaus SBL	0.13%	0.0008%
Spamhaus PBL	5.22%	0.0001%
Spamhaus XBL	1.38%	0.0023%

Table 2: Average intersections of blacklists with whitelist

lists entire net ranges that Internet providers declare as end-customer dialin IPv4 address space. Many mail receivers rely on it for blocking SMTP connections, which makes its use equal to a blacklist. For that reason we included the Spamhaus PBL in our measurements. As the intersection table shows, this coarse-graind listing procedure leads to many cross-listings. Due to the big IPv4 adress range covered by the Spamhaus PBL, however, only a small fraction of entries was listed in the whitelist.

These absolute numbers of cross-listed IP addresses are a good starting point to evaluate the impact of cross-listings. However, they completely ignore the relevance of an IP address in terms of mail volume. High volume mail senders using a cross-listed IP address have a larger practical impact than low-volume senders. Thus, we expand our model in the following subsection.

2.3 Mail volume of cross-listings

Despite the cross-listed IP addresses, our next model takes into account the mail volume originating from each sender. Although we cannot measure the mail volume per sender, we use a methodology to estimate which ratio of the global mail traffic is sent by cross-listed mail servers. We use DNS request data gathered at a mirror of the NiX Spam blacklist to estimate which ratio of mail originated in cross-listed IP addresses. In other words, we count the number of blacklist DNS queries that requested a cross-listed IP address and compare this number with the total number of requests.

At our blacklist mirror, we received 18 mio. DNS queries per day on average. Each request was likely issued by a mail server that looks up the IP address of a mail sender at our blacklist mirror. Of these requests, on average 2,855 requests per day queried a cross-listed IP address. In other words, every 6,300th SMTP connection (0.016%) is initiated by a cross-listed mail sender. Whereas these numbers measured at our blacklist mirror may sound low, estimations on the global mail volume reveal a big impact. As a consequence, we think it is worthwhile to evaluate in detail, whether black- and whitelists can be improved by rechecking the set of cross-listed IP addresses. In the next section, we use the information of cross-listings to improve DNS based anti-spam measures.

3. CROSS-LISTING CLASSIFICATION

IP addresses that are part of the gray IP area are candidates of wrong listings in a DNS based list. As discussed earlier, however, there may also be reasons behind a cross-listing. In this section, we will present a methodology that reveals conflicts between black- and whitelists. We will describe a classification scheme that allows to identify cross-listings as list error candidates. These candidates can be inspected manually and, if applicable, should be removed from the black- or whitelist.

Our classification scheme is based on the following measures:

- **Cross-listing constellation**:
  Our metrics consider the exact intersection constellation, i.e. how many lists of which type list a given address. Some lists are considered more valuable and trustworthy than others and are weighted accordingly. We base our weighting on expert advice.

- **Whitelist classification**:
  Whitelists often use classification schemes for their entries that indicate what kind of mail server is behind an entry. Some whitelists even specify which spam ratio can be expected from a whitelisted source. We integrate this classification in our scheme.

- **Reverse DNS query**:
  We try to resolve the domain name behind the cross-listed IP address. If the domain name is not resolvable, no well-configured mail server can be expected at this address. If a domain name is returned, we check whether it indicates that the mail server is legitimate or operating on a dialup address, respectively. In addition, we check whether the IP address that is behind the resolved domain name shares the same network (/24, /16) with the cross-listed IP address. This helps to possibly identify domain name fakes.

We use these criteria and order all cross-listed IP addresses by two schemes. Whereas the first scheme tries to identify wrong listings on a whitelist, the second ordering indicates possible legitimate mail servers on blacklists. The following two subsections explain the schemes used to derive this order of IP addresses.

3.1 Detecting blacklist deficiencies

Although DNS blacklists have been applied effectively as anti-spam measure over the last years [3], they bring along some risks. If used to block entire SMTP connections, blacklists prevent any incoming mail sent from a mailer using a listed IP address. This is particularly bad if mail servers largely send legitimate mail, but were listed on a blacklist because a small fraction of the mail volume was unsolicited. Our approach is to find these mailers and order them by an estimated level of spam that originates from them.

In our classification, the following criteria lead to a no-spammer-score:

- The quality of blacklists that list the probed IP address is lower than the quality of the whitelists.

- The spam level as classified in the whitelist, if applicable, is low.

- The resolved domain name behind the probed IP address has the pattern of a mail server (e.g. *mail*, *mx*, *smtp*)

- The no-spammer-score is decreased, if the resolved domain name cannot be verified. This happens if the IP address behind the resolved domain name is in a different /16- or /24-network than the probed IP address.

3.2 Detecting whitelist deficiencies

DNS whitelists help to minimize the amount of legitimate mails that is erroneously rejected by anti-spam measures, in particular blacklists. Usually, a whitelisted sender is guaranteed to get his mail passed through all anti-spam filters applied at the receiver's site. Whereas this is practical and prevents from false positives, the set of whitelisted senders should be chosen carefully. Supposedly legitimate and listed senders may send spam and, intentionally or not, abuse their good reputation. Therefore, our approach tries to indicate from which whitelisted IP addresses spam is originating.

In our classification, the following criteria lead to a spammer-score:

- The quality of blacklists that list the probed IP address is higher than the quality of the whitelists. If multiple blacklists share the probed IP address, the score is raised accordingly.
- The spam level as classified in the whitelist, if applicable, is considerably high.
- The resolved domain name behind the probed IP address has the pattern of an IP address used by end-customers (e.g. *ppoe*, *dialin*, *ppp*).
- The reverse DNS resolution of the probed IP failed, the resolved domain name was incorrect (e.g. localhost) or the domain name verification failed.

3.3 Score computation

In section 3, we discussed a method to classify whether a cross-listing is a likely spammer or a legitimate mail sender. In this subsection, we will summarize our classification results. For the analysis, we focus on cross-listings of the whitelist with maxium 1% spam ratio. Recall that we consider an IP address as listed in this whitelist as soon as it managed to have a maximum spam ratio of 1% for at least 1 out of the 20 days measurement period.

For these cross-listings, we computed a score that is assigned to each IP address. A high score indicates that there is a likely legitimate host behind the IP address, whereas a low score is assigned to likely spammers. In detail, every IP address has a base score of 0. We then computed the score as follows:

- Add 10 or substract 10 from the score, if the resolved domain name shows either a mail server pattern or dialup user pattern, respectively.
- Substract 5 from the score, if the PTR record of the IP address was not resolvable.
- Substract 10 from the score, if the forward-confirmed reverse DNS resolution failed.
- Add 10 to the score, if the cross-listed IP address was found in any other whitelist.
- Substract x from the score, if the cross-listed IP address was found in any other blacklist. x is a well-chosen value between 1 and 10 and depends on the quality that we assumed of a blacklist. For example, $x = 9$ for all Spamhaus.org lists, other blacklists having $x = 3$ or lower.

4. EVALUATION

In this section, we evaluate our classification method. First, we show that the DNS backwards resolution reveals many suspicious hostnames. We will then evaluate the overall classification score.

4.1 Domain name resolution

As discussed in section 3, we resolved the domain name that was behind the IP address of a cross-listing. We summarize our results in Table 3. For the DNS reverse resolution, we queried cross-listings of three different whitelists. Each of the whitelists contains legitimate hosts only. However, the maximum spam ratio threshold is different for each whitelist. Whereas the first list contains senders that have a spam ratio of maximum 1%, the third list even contains "legitimate" senders up to a spam ratio of 90%. We will include a discussion in our conclusion about this conflict. For now, given the different whitelists, we devided the results into three groups (columns in Table 3).

spam level	*up to 1%*	*up to 10%*	*up to 90%*
total IPs:	2,092	2,734	20,474
mx:	36	44	216
smtp:	73	97	141
mta:	37	49	62
static:	68	90	681
outbound:	3	4	5
mail:	220	260	476
SUM:	437	544	1581
NXDOMAIN:	546	727	4,720
localhost:	6	10	188
dial:	114	128	317
dyn:	197	259	2,415
proxy:	3	3	7
ppp:	17	34	963
ppoe:	4	10	660
SUM:	335	434	4362
out of /24:	32	42	377
out of /16:	27	34	350

Table 3: DNS resolution results

In this section, we will focus on the whitelist least prone to spam. That is, each entry had a spam ratio of maximum 1% for at least one day during the test period (20 days). We observed 2,092 cross-listings based on this whitelist. Of these 2,092 IP addresses, 1,546 DNS names could be resolved. Very much to our surprise, the PTR records of the remaining 26.1% of cross-listings were not resolvable (NXDOMAIN). We suspect that these senders will typically have issues sending mail, because many mail servers are configured to perform a DNS reverse lookup before accepting SMTP connections.

Of the resolved domain names, 437 names pointed to typical legitimate mail senders. 335 domain names had the pattern of typical dialin-addresses. The remaining cross-listings did not follow any pattern that we consider as typical neither for mail servers nor for home users. In 32 cases, when resolving the IP address behind the resolved PTR record, this IP address was in a different /24-network than the original IP address. In most of these cases (27), the IP addresses were even in different /16-networks.

To summarize, the DNS resolution revealed that IP addresses of more than a fourth of all cross-listings were not resolvable. More than a fifth of the resolved domain names showed a pattern of home users, where one would typically not expect legitimate mail servers. Finally, when we realized that somes servers even responded with "localhost" as PTR record, we decided to go one step further to check the quality of the cross-listings.

4.2 Classification score

The score was computed for all of the cross-listings that had a mail volume of more than 10 mails during the measurement period. The mail volume data was provided by a large anti-spam operator. Although this is certainly not a perfect snapshot of the global mail volume, it is sufficient to filter very low-volume senders. Of the 2,092 cross-listings, we computed the score for 1,681 IP addresses.

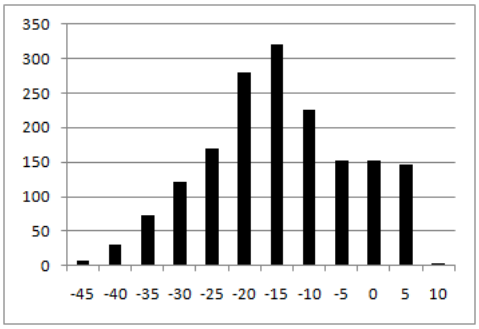

Figure 2: Score-distribution of cross-listings

Figure 3 shows the corresponding score distribution. Most of the cross-listings have a low score, indicating that spam is sent from these IP addresses. 301 IP addresses had a positive or at least neutral score.

To evaluate our classification mechanism, we correlated our results with the ratio of clean mail originating in each of the cross-listed IP addresses. In addition to the mail volume, the anti-spam provider also provided percentages of legitimate mail sent by an IP address. With this data as ground truth, it was possible to assess whether we developed a well-working classification mechanism.

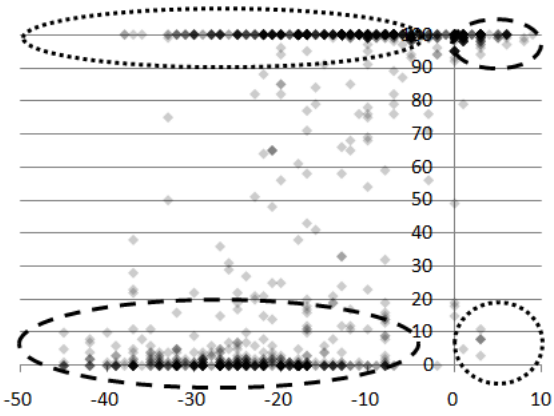

Figure 3: Correlation between score and ratio of legitimate mail

Figure 3 shows the distribution of all cross-listings, scaled on the classification score (x-axis) and the percentage of legitimate mail sent by that IP address (y-axis). Most of the cross-listed IP addresses had either a very low ($< 5\%$), or a very high ratio of legitimate mails ($> 90\%$). Each scattered point is semi-transparent, such that many values at a single point appear bolder.

We devided Figure 3 into four sections. The sections that are marked with dashed ellipses show cases where our score correctly identified whether hosts send spam or not. In the lower-left area, the low score correctly indicates that only a small fraction of mail of these cross-listings is legitimate. In the upper-right area, an at least neutral score correctly classified legitimate senders. However, in other cases our classification mechanism failed. In the lower-right area, some IP addresses are identified as legitimate senders, although they have a fairly small ratio of legitimate mail. Moreover, in the upper-left area of the graph, we see that many cross-listings were classified as spammers although they vastly send legitimate mail.

cat.	*legitimate*	*score boundary*	*no. of cross-listings*
1	$< 20\%$	≥ 0	9
2	$< 20\%$	< -5	578
3	$> 90\%$	≥ 0	271
4	$> 90\%$	< -5	680

Table 4: Summarized classification results

Table 4 summarizes the four distinctive classification categories. Overall, 849 servers were correctly classified as source for spam or solicited messages, whereas our mechanism gave a wrong reputation to 689 IP addresses. Unfortunately, based on a classification error of 45%, we have to admit that our automated classification of the gray IP area works only slightly better than simply guessing. Moreover, 143 IP addresses had a medium ratio of legitimate mail that was not sufficient to call it either a spammer or a legitimate host.

However, having this negative result, we are positive about conclusions that we can draw from this evaluation. On the one hand, we are convinced that our approach of classifying IP addresses (as discussed in section 3) is a best-possible method taking into account a wide set of features that can be attributed to IP addresses. On the other hand, we show that the derived score is not a good indicator of the likeliness of receiving spam from a given IP address. At this point we stress the fact that we did not modify the weights that we gave to construct the score to obtain a better classification result. Although we think this would be possible and the classifier could be further optimized, no matter the improvements, it would never reach a satisfying level with a sufficient low error rate. In hindsight, when recalling our reasoning of cross-listing in section 2, basing a classification method on blacklists was made to fail. However, we learned lessons from our work, as we will summarize in the conclusion section.

5. DISCUSSION AND FUTURE WORK

We presented a mechanism to classify cross-listings towards a spammer likeliness based on DNS name resolution and combining listing results from multiple weighted blacklists. Unfortunately, this scheme failed so that we do not see any future work that can improve the mechanism. We

even go further and argue that it is impossible to automatically classify cross-listings as good or bad, simply because there is a perfect reason for them to exist. Although we are pessimistic about improvements of our classification mechanism, there are some research areas that we will explore in the future.

In this work, we highlighted trends that have been changing the use of IP blacklists. Knowing these trends, we think that blacklists will be increasingly prone to false positives in the future. Some work has been done to examine false positive rates of certain blacklists individually in the past [4]. However, to the best of our knowledge, there is no large-scale approach to measure the false positive rate of John Doe's mail account. We will investigate approaches that can measure the false positive rate of blacklists and observe whether we can see influences of the mentioned trends.

Finally, a very interesting trend is the introduction of IPv6. On the one hand, some people argue that blacklists will not be capable of handling the huge address space (2^{128} instead of 2^{32}). Indeed, this may become a problem if too many IP addresses will be assigned over time. On the other hand, if fixed IP addresses are assigned to end-users in the era of IPv6, then the problem of listing dynamic IP address pools will be mitigated. We will observe this trend and expand our various research activities in the are of anti-spam to IPv6.

6. CONCLUSIONS

This paper attempted to substantiate and quantify the increased occurrence of false positive messages that are expected when using a combination of blacklists and whitelists. To validate (or refute) this supposition we attempted to create an analysis system which correlated the incidence of spam sending to the white or blacklist IP classification. However, after establishing such a methodology to evaluate blacklists and whitelists we were unable to create a classification system that accurately determined the veracity of IP entries listed on white/black lists. Nonetheless, the failed attempt at classification has led to several verifiable findings:

- The analysis clearly demonstrates that many entries are incorrect for both the black and white lists we evaluated. These misclassifications provide the ground for legitimate mail to become false positives. Given our imperfect whitelists, already every 6,300th SMTP connection is established from a cross-listed IP address. However, this lower bound for the gray area indicates the minimal magnitude in which the use of IP lists leads to false positives.
- Many of the IP listings fall into grey areas of which no reproducible method of analysis could accurately recommend IP address removal from the black- or whitelist and for whom the behavior varies greatly from legitimate to spam.
- Many of the current spammer techniques directly undermine the usefulness and thereby, accuracy, of black- and whitelists. This inaccuracy is only compounded by the false positive rates already prevalent in many content filtering solutions.

As demonstrated by this research, despite the prevalent use of blacklists, they must be augmented by other anti-spam techniques (e.g., content filtering). Our research clearly indicates that the assertion of increased occurrences of false positives when using blacklists is not only sound but certainly quantitatively verifiable. The next iteration of this research will focus on measuring this false positive rate on an average user's email account.

7. REFERENCES

[1] R. Clayton. Stopping Spam by Extrusion Detection. In *Conference on Email and Anti-Spam (CEAS)*, 2004.

[2] R. Clayton. Stopping Outgoing Spam by Examining Incoming Server Logs. In *Conference on Email and Anti-Spam (CEAS)*, 2005.

[3] European Network and Information Security Agency. Provider Security Measures - Survey on Security and Anti-Spam Measures of Electronic Communication Service Providers. In *Deliverable 2.1.6 of ENISA's Work Programme*, 2007.

[4] A. Iverson. DNSBL Resource: Statistics Center - http://stats.dnsbl.com/, 2009.

[5] J. P. John, A. Moshchuk, S. D. Gribble, and A. Krishnamurthy. Studying Spamming Botnets Using Botlab. In *Usenix Symposium on Networked Systems Design and Implementation (NSDI)*, 2009.

[6] C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. Voelker, V. Paxson, and S. Savage. Spamalytics: An Empirical Analysis of Spam Marketing Conversion. In *Proceedings of ACM Conference on Computer and Communications Security (CCS)*, 2008.

[7] M. Kokkodis and M. Faloutsos. Spamming Botnets: Are we losing the war? In *Conference on Email and Anti-Spam (CEAS)*, 2009.

[8] Lashback. LashBack - The Email Compliance Authority - http://lashback.com/, 2009.

[9] Messaging Anti-Abuse Working Group (MAAWG). Message Sender Reputation Concepts and Common Practices. In *White Paper*, 2009.

[10] G. Mori and J. Malik. Recognizing Objects in Adversarial Clutter: Breaking a Visual CAPTCHA. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2003.

[11] A. Ramachandran, D. Dagon, and N. Feamster. Can DNS-Based Blacklists Keep Up with Bots? In *Conference on Email and Anti-Spam (CEAS)*, 2006.

[12] A. Ramachandran, N. Feamster, and S. Vempala. Filtering Spam with Behavioral Blacklisting. In *Proceedings of ACM Conference on Computer and Communications Security (CCS)*, 2007.

[13] Spamhaus. The Spamhaus Project - http://spamhaus.org/, 2009.

[14] B. Stock, M. Engelberth, F. C. Freiling, and T. Holz. Walowdac Analysis of a Peer-to-Peer Botnet. In *European Conference on Computer Network Defense (EC2ND)*, 2009.

[15] B. Stone-Gross, M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel, and G. Vigna. Your Botnet is My Botnet: Analysis of a Botnet Takeover. In *ACM Conference on Computer and Communications Security (CCS)*, 2009.