

# Developing Security Reputation Metrics for Hosting Providers

(Step Right Up! Metrics Galore! )

A. Noroozian      M. Korczyński      S. TajalizadehKhoob      M. van Eeten  
Delft University of Technology

## Abstract

Research into cybercrime often points to concentrations of abuse at certain hosting providers. The implication is that these providers are worse in terms of security; some are considered ‘bad’ or even ‘bullet proof’.

Remarkably little work exists on systematically comparing the security performance of providers. Existing metrics typically count instances of abuse and sometimes normalize these counts by taking into account the advertised address space of the provider. None of these attempts have worked through the serious methodological challenges that plague metric design.

In this paper we present a systematic approach for metrics development and identify the main challenges: (i) identification of providers, (ii) abuse data coverage and quality, (iii) normalization, (iv) aggregation and (v) metric interpretation. We describe a pragmatic approach to deal with these challenges. In the process, we answer an urgent question posed to us by the Dutch police: ‘which are the worst providers in our jurisdiction?’. Notwithstanding their limitations, there is a clear need for security metrics for hosting providers in the fight against cybercrime.

## 1 Introduction

Hosting providers are companies that provide servers via which customers can make content or services available on the Internet *e.g.* websites, email or support for multiplayer gaming. As with virtually all services on the Internet, they are abused for criminal purposes as well. A wealth of research has identified how hosting infrastructure shows up in various criminal business models. Think of phishing sites, command-and-control servers for botnets, child pornography, malware distribution, and spam servers [1].

Nobody contests that hosting providers play a key role in fighting cybercrime. Much of the criminal activity runs on compromised servers of legitimate customers, some on servers rented by the criminals themselves. In either case, the hosting providers typically becomes aware of the problem only after being notified of the abuse. Their response to abuse reports varies widely,

ranging from vigilant to slow to negligent or even bullet-proof [1, 2]. To empirically measure which of these responses is actually occurring has proven to be very challenging. Existing metrics of hosting provider security typically count instances of abuse within an Autonomous System, sometimes normalized by the size of the advertised address space [1, 3, 4]. None of these attempts adequately account for the serious methodological challenges plaguing such metrics.

In this paper, we present a systematic approach for developing metrics for hosting providers. It enables us to identify and discuss the main challenges: (i) identification of providers, (ii) abuse data coverage and quality, (iii) normalization, (iv) aggregation and (v) metric interpretation in light of the heterogeneity of hosting providers. Additionally we present a pragmatic approach to deal with these issues.

This study is part of an ongoing collaboration with the Dutch National High Tech Crime Police, the Authority for Consumers and Markets, the Public Prosecutor and the Dutch Hosting Provider Association. The objective is to answer an urgent question posed by the police: which are the worst providers in our jurisdiction?

The question illustrates there is a clear need for security metrics for hosting providers, notwithstanding their limitations. Reducing cybercrime is as much a problem of incentives as it is a technical issue [5]. Without reliable signals we cannot tell which provider is vigilant, lax, negligent or outright criminal and it will be very difficult to move the sector towards more secure practices. *Information asymmetry* erodes the incentives of providers to invest in security. Reliable metrics can (i) signal security performance to customers, upstream and downstream providers, law enforcement and other stakeholders (ii) enable benchmarking of providers, and (iii) help identify the effectiveness of security practices and policies.

The main contributions of this paper are as follows: (i) we systematically outline the process to develop security reputation metrics for hosting providers, as well as the methodological challenges encountered along the way, (ii) we improve existing techniques for mapping abuse to hosting providers and for taking into account the size

of hosting providers in computing reputation scores, (iii) we present a pragmatic approach to produce metrics for the Dutch market, developed in collaboration with key stakeholders.

## 2 Background

Hosting providers come in many shapes and sizes and offer portfolios of services: from relatively expensive dedicated physical machines to virtual private servers (VPS) to the cheaper options of shared hosting or even so-called free hosting. In each service, the role of the provider vis a vis the customer is different. On a dedicated machine, and to a lesser extent on a VPS, the customer controls the entire software stack, whereas on shared hosting, many customers operate under restricted privileges on a machine they share with many other users. Free hosting services limit user control to the extreme.

Depending on the type of customer, hosting providers play a different role in protecting their customers against compromise by patching servers, cleaning and monitoring for abuse. Similarly, providers need to protect the rest of the Internet against potentially malicious customers by putting in place different checks and restrictions which depends on the service contract with that customer.

Next to the occurrence rate of abuse, the uptime of abuse also reflects hosting provider security practices. On one end of the spectrum, vigilant hosting providers remove malicious content often within hours of its discovery, in the middle there are some providers that respond more slowly and more selectively and on the other extreme are the so called ‘bulletproof’ hosting providers that seem to ignore all abuse notifications.

There has been a lot of speculation over the incentives of providers. A shared hosting provider, for example, could act against abuse more directly because its customers have only limited control over the machines that they use. On the other hand, shared hosting is a highly competitive market with low margins, so investing in security is not likely to be a high priority. The only way forward is to replace speculation with reliable empirical evidence of abuse rates across providers.

## 3 Overview of Approach

Our approach for calculating reputation metrics is partly guided by the goal to allow our collaborators to engage in meaningful discussions based on reliable empirical techniques. To this end, we produce two types of security indicators for hosting providers based on data available in public and private abuse feeds: (i) *Occurrence of abuse*: an indicator based on counting occurrence of abuse, and (ii) *Persistence of abuse*: an indicator based on how long the abuse was present. Distinguishing between occurrence of abuse and the response of a hosting provider to

abuse as independent measures of performance is important. While the occurrence of abuse is to some extent inevitable due to technical vulnerabilities and related to organization size and attacker characteristics, persistence of abuse indicates attitude towards dealing with abuse and mainly relates to defender characteristics. In conjunction, these independent indicators provide a better understanding of the overall security performance of a hosting provider.

Figure 1 illustrates a high level overview of the complete procedure to produce these indicators. Here, boxes represent inputs/outputs to each step while arrows transformation steps on data required for each step. The process is quite generic and outlines the steps that any reputation metric requires to arrive at final scores. In executing these steps there are challenges that need to be overcome and choices that have to be made that will undoubtedly effect the reliability and interpretation of the reputation metric. In what follows we systematically walk the reader through the steps of the process meanwhile highlighting challenges related to each step and the possible effects on the overall metric and its interpretation. A more detailed analysis of some choices and their effect are presented later in Section 10.

## 4 Step 1 - Abuse Mapping

Identifying hosting providers is not straight forward since they do not directly map onto entities with which underlying Internet protocols work or what abuse data capture. The first decision that needs to be made is to *identify what a hosting provider is*.

### 4.1 Identifying Hosting Providers

To produce reputation metrics for Dutch hosting providers, we made the (common [1, 3, 6, 7]) assumption that hosting providers will have an associated Autonomous System Number (ASN). Consequently, we considered any AS which routed IP addresses that geo-located to the Netherlands as a Dutch hosting provider.

While the assumption may hold in general, ASes could refer to Internet Service Providers (ISP), Internet exchange points, banks, governmental institutions, universities and in general non-hosting entities. Without a deeper analysis of the ASes, such an assumption may lead to considerable error in mapping abuse onto hosting. Even when an AS does refer to a hosting provider, complexity still exists. Providers may have multiple ASNs or there might be multiple organizations which own a part of the IP space in an AS or *resellers* who lease infrastructure from the AS. Some ASes also route traffic destined for IPs owned by *peers*. Certain legitimate services (e.g. CloudFlare) may act as proxies and hide providers. As a result abuse associated with small organizations with registered IPs in ASes may end up attributed to the AS

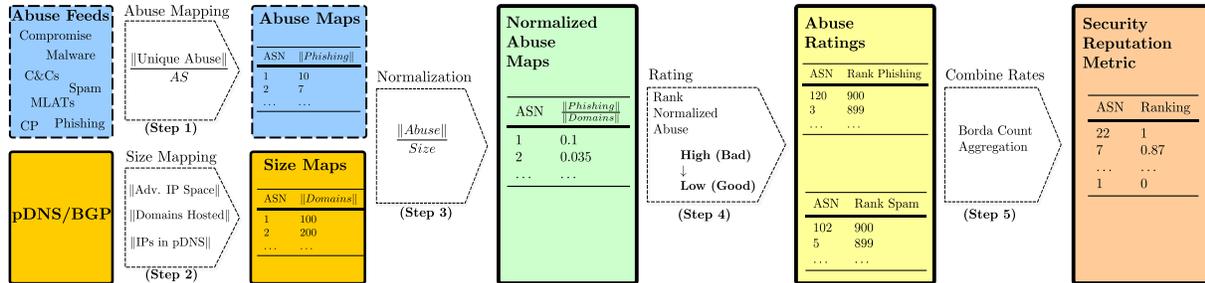


Figure 1: Reputation metric development process

from which the infrastructure is leased. We have opted to analyze abuse for organizations with registered IPs inside ASes in future work.

One method to better identify hosting providers and identifying organizations under each AS, is to analyze IP ‘ownership’ using Maxmind’s GeoIP ISP Database. Utilizing such information results in a more fine grained mapping which mitigates the mapping problems discussed above. Nevertheless, this approach has complications of its own such as non-standardized WHOIS data formats where the same organization might appear with multiple names that are non-trivial to relate to each other. For example, the Dutch provider Leaseweb might appear under any of the following additional names: Leaseweb Asia Pacific. ltd., leaseweb1.iomadserve.com.

## 4.2 Unit of Abuse

The second key decision is about the *unit of abuse* or *how to count the abuse data*. Unlike other hosting metrics which typically count distinct IP addresses as the unit for abuse [1, 3, 7], our approach considers unique 2nd-level domain-IP pairs -  $\langle 2LD, IP \rangle$  - as the unit of abuse. From this point on, we use the terms ‘2LD’ and ‘domain’ interchangeably, unless the context requires otherwise. Simply counting the number of abusive IP addresses largely underestimates abuse from shared hosting services since criminals may use the same IPs for various purposes. For example a compromised server may host a phishing website and also be used for spreading malware. Furthermore, the number of domains is a better proxy for the number of customers of the provider, which is valuable to include in approximating its size. Last, this definition also maximizes the value of our feeds as measured by their differential contribution [8].

Counting pairs of  $\langle 2LD, IP \rangle$  mitigates the problem but is not perfect. In some cases it is appropriate to count  $\langle FQDN, IP \rangle$  pairs (e.g. malicious domain generation algorithms), or even  $\langle URL, IP \rangle$  pairs (e.g. child abuse content concentrated under the same domain with varying paths in the URL).

## 4.3 Data feeds

A separate decision in mapping abuse is *what data feeds to use*. A wide range of abuse on the Internet is associated with hosting. Hosts are used as malware drop zones and to host phishing pages designed to steal sensitive information. Botnet command and control (C&C) servers are also hosted [7]. Other types of hosting related abuse includes child pornography, SEO schemes, spam and counterfeit goods stores. Not all criminal activity can be observed in a way that can be attributed to the infrastructure of a specific hosting provider. Think of hidden services on TOR. Even if it can be observed, the criminal activity might not be captured in abuse data feeds, which are often produced by automated means. This implies that abuse feeds are always partial and of varying quality. This is a well known fact [8, 9]. Needless to say, criminal activity that is not captured in the abuse data included in a metric, forms a blind spot of that metric. This suggests to include as broad a spectrum of abuse feeds as possible.

We collect a range of feeds and blacklists from private, public, commercial, and governmental sources. Table 1 gives an overview of these data feeds. The data spans over the entire duration of 2014 (with the exception of the SHC and SHS which span over the 2nd half of 2014). The majority of our feeds, do not share much information on the exact collection methodology. We did not include some of the available spam feeds because our analysis of the data revealed these to be mostly related to compromised end users of ISPs rather than hosting companies.

Table 1: Data Feeds and Statistics

Abuse Type	Feed	Organization	Samples			
			$\langle Domain, IP \rangle$		IPs	
			Total	Excl.	Total	Excl.
Malicious Hosts	SHC	Shadowserver	3957	3615	2260	1321
Malicious Hosts	SHS	Shadowserver	7632	7489	1100	816
Malware	SBW	StopBadware	15204	14757	7702	6170
Botnet C&Cs	ZEUS	Abuse.ch	50	27	72	35
Phishing	PHISH	Phishtank	2278	1780	1377	
Phishing	APWG	APWG	3060	2430	1886	1101
Take Down Request	MLAT	Dutch Police	1347	1202	1433	1202
Child Pornography	MELD	Meldpunt	725	584	417	242
Total			34253	31884	16247	11491

In general, data quality relates mainly (but not only) to: (i) coverage (*What is the overlap between the different feeds?*), (ii) purity (*How much of the blacklisted domains truly host malicious content?*) However, it is not always possible to assess the coverage or purity of a feed since mainly its details are not well documented [8].

**Coverage.** Previous overlap analysis of blacklists that cover different types of abuse concludes that - although existent - there is little overlap in terms of the abuse associated with each ASN [9]. We reach similar results especially when  $\langle 2LD, IP \rangle$  pairs are the unit of abuse.

Clearly the feeds differ substantially in terms of the volume of reported abuse samples. For example, the professionally oriented *SBW* feed contributed over 15,000 samples, while the non-profit *ZEUS* feed three orders of magnitude less domain-IP pairs. In terms of the total number of IPs, *SBW* reports almost two times less unique IPs than distinct  $\langle 2LD, IP \rangle$  pairs whereas the *ZEUS* feed reports more IPs because some Zeus config, binary, and drop zones are hosted solely on IP addresses. Moreover, the differences between  $\langle 2LD, IP \rangle$  pairs and IPs indicate that many domains used for criminal activity are mapped to a smaller number of IP addresses which could be the result of shared hosting services. Across our feeds, 93% of all  $\langle 2LD, IP \rangle$  pairs and 71% of all IPs for *all domains* were exclusive to a single feed (cf. *Excl.* column in Table 1). We refer to samples as *exclusive* when they appear only in one feed.

Figure 2 illustrates pairwise feed intersections as a matrix, with unique  $\langle 2LD, IP \rangle$  (left) and unique IPs (right) as the unit for abuse respectively. Here darker shades of grey represent higher overlaps. For instance, in Figure 2 (left), the overlap between *MLAT* and *MELD* indicates 124  $\langle 2LD, IP \rangle$  pairs in common. This overlap constitutes 9% of the *MLAT* feed. In comparison, 124  $\langle 2LD, IP \rangle$  pairs represents 17% of the *MELD* feed. The rightmost column indicates the absolute number and the percentage of samples that the feed has in common with all other feeds combined. The amount of overlap in the rightmost columns confirms the simultaneous use of IPs for different malicious purposes. This also further supports our choice of abuse unit.

Finally, the relatively small overlap among our chosen data feeds in terms of  $\langle 2LD, IP \rangle$  indicates the suitability of these feeds, still, other feed characteristics need to be analyzed to further establish suitability.

**Purity.** All abuse feeds contain false positives. The main question is which samples should be considered false positives and excluded and which should not? We define false positives as websites maintained by legitimate users that do not serve any malicious content and are incorrectly blacklisted.

Some domains are legitimate but point to servers that host malicious content. For example, we consider URL

shortening services such as `goo.gl` or `bit.ly` as false positives. However, other legitimate websites such as free web hosting providers (e.g. Hostinger), or cloud storage services (e.g. Imagezilla.net or Dropbox) are misused by criminals and included in our analysis.

Moreover, a certain portion of abuse feeds include benign domains. We analyze benign domains that appear in the Alexa top 25 thousand domain list to evaluate the prevalence of false positives. Although we do not provide a real-time verification of blacklisted Alexa domains, we perform *a posteriori* analysis to further establish the suitability of our chosen abuse feeds.

Due to space limitation we only briefly discuss the analysis and do not include the details. Overall we find a low number of Alexa domains in the abuse feeds. Nevertheless, there are major differences among the types of Alexa domains per feed. For example, through manual analysis of a random sample from the *SHS* feed we found that approximately 30% of this feed’s Alexa domains were file sharing services most probably used to host malicious content and thus relevant to include. On the other hand, we also find some examples of popular websites like `msn.com`, or `microsoft.com` that are presumably used by compromised machines to test network connectivity. In the case of the *PHISH* feed we find that a significant number of ranked domains that are either false positives (e.g. banks and other legitimate services), or are not appropriate for the type of analysis that we provide (e.g. URL shorteners). The majority of the ranked domains for both *MLAT* and *MELD* represent file and adult content sharing services. As systematic false positives, or unrelated web services in our feeds do not constitute a large number we therefore have opted to include such ranked domains in our analyses.

## 5 Step 2 - Size Mapping

Reliable reputation metrics need to account for a commonly observed trend that larger providers also experience a larger amount of abuse.

A common yardstick for measuring the ‘size’ of hosting providers is the number of IP addresses routed by its corresponding AS in the BGP protocol [1, 3]. Nevertheless not all IPs routed by an AS are used for hosting content nor are they directly in use by the AS. IPs may be leased and used for other purposes. Inaccuracies in size estimation may negatively impact the reliability of a metric in that they can lead to misleading results. Nevertheless, due to simplicity of calculating, advertised IP space remains an attractive choice for size estimation.

We propose (and use) two additional size estimators: (i) the *number of hosted 2LD* and (ii) the *number of IP addresses used to host content* per hosting provider. To calculate these estimators we use the historical passive DNS (pDNS) data provided to us by Farsight Security.

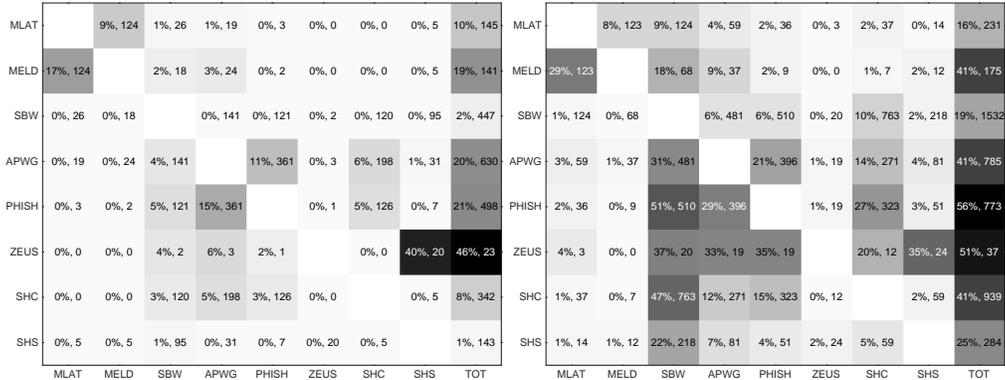


Figure 2: Pairwise overlap of feeds with unique  $\langle 2LD, IP \rangle$  (left) and IPs (right) as unit of abuse

This data records domain name resolution queries collected over the entire duration of 2014 which we use to count the number of unique 2LDs and their matching IP addresses. These counts are subsequently mapped to ASes routing the IP addresses and used as an estimation of the size of the provider. Here the quality of the resulting estimates is highly dependent on the coverage of the pDNS data. As long as the pDNS data has a reasonable coverage of all registered 2LDs, it can be used to produce reasonable size estimates. We have crosschecked the number of unique 2LDs observed in the pDNS data with the number of 2LDs of gTLDs and ccTLDs present in zone files of new gTLDs that we obtained under agreement from ICANN<sup>1</sup> in addition to ccTLD sizes reported by APWG<sup>2</sup> at the end of 2014. Extrapolating from these results we have concluded pDNS to be a reasonably reliable source to estimate hosting provider size.

There are potentially other conceivable estimators for hosting provider size such as the number of customers. Nevertheless, the scarcity of data to base such estimates on is a largely limiting factor in this respect.

## 6 Step 3 - Normalization of Abuse

Given the output of abuse mapping and size mapping, the next step in the metric production process is to normalize abuse by a size estimate. This leads to  $S \times N$  normalized abuse mappings where  $S$  is the number of size maps produced earlier and  $N$  the total number of abuse maps corresponding to analyzed blacklists. A key question here relates to *interpretations that can already be made from normalized abuse data*.

All size estimates have their advantages and disadvantages which have to be viewed as trade-offs. The most commonly used size estimator - routed IPs - is the easiest to calculate, but it suffers from systematically favoring large providers, since not all routed IPs are used for

hosting. Using the portion of the routed IP space that is used for hosting as the size estimator mitigates the problem, however, this is much more difficult to calculate. This estimate is also not free of systematic bias, because it favors hosting providers that have a disproportionately large amount of shared hosting. We can use the number of hosted 2LDs as the estimator, which would treat shared hosting fairly but would still underestimate the size of subdomain resellers and free-hosting services. The trend here is clear; normalized abuse has its blind spots, and needs to be taken into account especially for interpreting results at this stage.

It is important to note that some size estimates are more volatile than others due to the dynamic nature of the underlying processes. For example, the number of FQDNs hosted by a provider may change at a much faster rate than the number of 2LDs if an estimator based on FQDNs is used.

Normalized abuse, is already an indicator of security performance by itself. Note however, that normalized abuse is abuse type specific. For example, one can analyze normalized abuse based on the occurrence of malware on hosting providers and draw conclusions; however, this only provides a partial picture of the performance of hosting providers. Some providers might be much less strict about allowing malware spread from their servers than for example the hosting of child pornography [7]. In our case, we use all size estimators outlined in the previous section without committing to a specific one or considering one superior to others. The expectation is that the combination of these can overcome the deficiencies of each. This matter is further explored in Section 10.

Finally, note that when talking about metrics based on the up-time of abuse, size corrections are not appropriate. In such cases it is common to use mean or median up-times instead of normalized abuse.

<sup>1</sup><http://newgtlds.icann.org>

<sup>2</sup><http://docs.apwg.org>

## 7 Step 4 - Rating of Abuse

Given the normalized abuse maps, the next step in the process, calculates rankings over all maps to produce rankings. Rankings are one way of unifying the scales on which normalized abuse is measured and allows cross comparisons over categories of abuse. For example, comparing the security performance of a hosting provider in terms of how well it manages to mitigate malware with its performance in terms of how well it mitigates phishing is not meaningful when based on normalized abuse. However, the comparison is meaningful over rankings.

Given the normalized abuse maps, our method for ranking hosting providers is as follows: We rank normalized abuse from high to low. This results in  $3 \times N$  rankings. The individual rankings may range between zero and  $R$ , the total number hosting providers. The worst rank,  $R$ , is assigned to the AS with the highest normalized abuse,  $R - 1$  to the second worst and so forth. ASes with equal normalized abuse are assigned equal ranks. If a normalized abuse map only contains data on for example 20 providers the ranking will range between  $[R - 20, R]$  with all providers for which no abuse was detected receiving the low rank of  $R - 20$ .

An important consideration in producing rankings is *information loss*. To illustrate this consider two hosting providers  $HP_1$  and  $HP_2$  that have a normalized abuse of 0.1 and 0.3 and have been assigned the ranks of 10 (worst performer) and 5 (5th worst) respectively. Our ranking is not *distance preserving* since the difference between the hosting provider ranks ( $10 - 5$ ) does not entail the same information as that of the normalized abuse ( $0.3 - 0.1$ ). That is, one unit of change in ranking could mean any number of changes in the unit of normalized abuse. As a result these distances cannot be interpreted in the same way. In ranking hosting providers, some information about the magnitude of the differences is unavoidably lost.

## 8 Step 5 - Aggregation of Rates

We now aggregate our rankings into one overall ranking that assigns scores in the range  $[0, 1]$ , where score 1 indicates the worst performer. The aggregation procedure considers every ranking as a voting preference over  $R$  candidates in an imaginary election. The election winner is effectively decided using a *Borda Count* vote aggregation method that basically counts how many times a certain candidate appeared in the 1st place, 2nd place, 3rd place (and so forth) in every ranking and decides the outcome based on all rankings.

An alternative approach could perform factor analysis and take into account the most contributing feeds when interpreting metric scores (cf. also Section 10).

We find, however, voting systems to be a useful analogy when thinking about aggregation. A useful aggregation method must have certain desirable properties, such as being intuitive. For example, if a particular hosting provider is the worst ranked performer in all categories of abuse, the reputation metric should reflect that by assigning the worst metric score to that provider and not to others. Certain methods of aggregation will not guarantee such properties and are therefore undesirable. We refer the reader to literature on different voting aggregation methods<sup>3</sup> for a better understanding of the properties of such methods and their limitations.

## 9 Step 6 - Metric Interpretation

Reputation scores need to be interpreted to guide policy and reduce information asymmetry around the security performance of hosting providers. However, correct interpretation of a metric without detailed knowledge of the various blind spots and biases of the process is difficult. Additionally the heterogeneity in the hosting provider landscape directly influences what conclusions can be drawn from the scores.

To illustrate the challenges of interpretation we briefly present some of our results here. Figure 3 plots the rankings of the 20 worst performers based on the occurrence of abuse. The plot demonstrates a large variance between the performance of providers that have comparable sizes. The results clearly indicate significant differences in how hosting providers deal with abuse. Here, the safest comparisons are among providers that have the most similar properties. As an example consider the two hosting providers colored in bright green, first the provider with the highest metric score and second the provider located approximately at  $(x \simeq 10^4, y \simeq 0.85)$ . These are very similar in all aspects and therefore it can be safely concluded that the provider with the lower score is performing significantly better than the worst performer due to its security policies and practices. To consider the worst provider as negligent or criminally engaged simply because it has the worst score is however a wrong conclusion to draw here.

Figure 4 compares the occurrence metric and the uptime metric of hosting providers. A cautionary note here is that our uptime metrics are based on only 2 data feeds from which up-times could be calculated. The weak relationship (Spearman's  $\rho = 0.38$ , Pearson  $r = 0.36$ ) between occurrence and up-time is expected as each captures a different aspect of hosting provider characteristics that relate to abuse (see Section 3). Clearly some providers experience large amounts of abuse while managing to quickly block the abuse (upper left region of the plot). Others, perform consistently bad in the sense that

<sup>3</sup>See e.g. <http://lorrie.cranor.org/pubs/diss/node4.html>

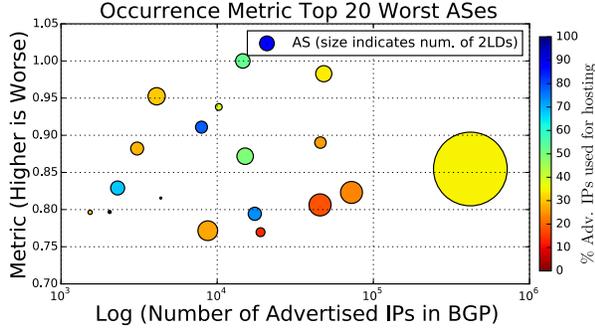


Figure 3: The 20 worst Dutch providers for abuse rate.

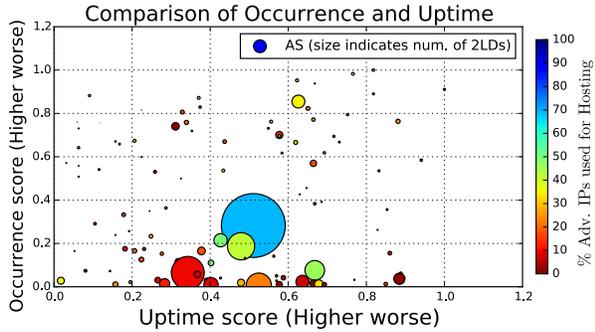


Figure 4: Comparison of occurrence and uptime metric.

they experience large abuse occurrence and are also slow to block it (upper right region of the plot). Nevertheless, we believe that the amount of occurring abuse and the response of a hosting provider to abuse are important aspects that need to be both measured separately to provide a thorough picture of security performance. Only now can we draw the conclusion that the worst performing hosting provider in terms of occurrence is probably negligent because it is also among the worst performers in terms of uptimes (see point with  $(x \simeq 0.8, y = 1)$  coordinates)

Finally when interpreting the results, one should also take the hosting provider business model into account. Hosting providers with a large portion of shared hosting customers have a larger role to play in cleaning up content than ones with dedicated hosting clients. It might very well be the case that the worst performer in terms of both occurrence of abuse and response to abuse provides solely unmanaged hosting to its customers and therefore not in the same position as its peers that provide mainly shared hosting. In this case the observed performance could simply be indicating the security of the hosting customers rather than that of the hosting provider itself.

## 10 Sensitivity Analysis

To better understand the impact of key design decisions, we undertake a brief sensitivity analysis of alternative specifications for (i) unit of abuse, (ii) abuse normalization, and (iii) metric aggregation strategies. We explore the robustness of the results by producing rankings based

on alternate methodological options for each of these decisions. We compare them to the ranking of the pragmatic approach we presented (called as benchmark ranking) by calculating the *Pearson correlation coefficient* among the top  $n = 100$  worst performing Dutch hosting providers.

**Unit of abuse.** We used unique  $\langle 2LD, IP \rangle$  pairs as the unit of abuse. We calculated an alternate ranking based on *unique IP counts*, the standard approach in the literature. The Pearson’s  $r$  among both rankings is 0,952. Perhaps more tangible: 16 ASes are in the top-20 of worst performers of both rankings. In other words, the metric is not very sensitive to either specification.

**Abuse normalization.** We calculated three alternate scores using the three size estimates for hosting providers: (i) the advertised IP space (ii) the advertised IP space that is used for hosting, and (iii) the number hosted 2LDs. The benchmark ranking used all three of them. The Pearson’s  $r$  for the alternative specifications to the benchmark ranking are 0,896, 0,909, and 0,6438, respectively. The results reveal a strong correlation between the IP space-based size estimators and the benchmark ranking, and a less strong correlation with the estimates based on 2LDs. Out of the top 20 worst performers in the benchmark ranking, only 7 ASNs were present in all alternate top 20 rankings. When comparing pairwise: the reference ranking shares 13 ASNs in common with those based on advertised IP addresses and hosted 2LDs, and 15 ASNs with the ranking based on IP addresses used for webhosting. In other words, using the number of hosted domains vs. estimates based on IP address space give significantly different results. By including all three size estimations, our metric specification mitigates that impact, while retaining the advantages of including domain name counts in abuse counting and size estimation.

**Metric aggregation.** We also compared our benchmark ranking with a ranking in which we assigned weights to each data source according to their comprehensiveness, i.e., the relative volume of each feed in terms of distinct number of exclusive  $\langle 2LD, IP \rangle$  pairs in the dataset (cf. Table 1). In summary, out of top 20 abused ASNs in the benchmark, 14 ASNs showed up in the weighted rankings and the Pearson’s  $r$  of the benchmark ranking and weighted ranking is equal to 0,791.

## 11 Related Work

Numerous studies have pointed to concentrations of abuse in certain networks, typically in the context of a specific criminal business model, e.g., spam [10], phishing [11] or malware [12]. Effects and policy implications of intervention at classes of intermediaries have also been studied in [10, 13]. The quality of abuse data has also been extensively covered [8, 9, 12]. [2] examines the

role of hosting providers in detecting abuse and reacting to user complaints for shared hosting providers. It paints a general picture that underlines the need for hosting reputation metrics. None of these studies try to develop reputation metrics from abuse data, however.

Closest to our work are [1, 3, 14]. [6] produces a weighted metric score for ASes. [1] includes only uptime data and focuses mainly on identifying the worst actors. We expand on this work by systematically addressing challenges not discussed there. These studies typically count IP addresses as the unit of abuse and use advertised IP address space as a normalization factor.

Industry attention to hosting has been along the lines of the Host Exploit Index (HE index) [4]. While valuable, their methodologies are not fully transparent and the parts that are, suffer from similar limitations as the academic work discussed above.

## 12 Conclusions

This paper has systematically worked through the many challenges of developing security reputation metrics for hosting providers. All conceivable metrics will suffer from various limitations, that much is clear. This is not to say that they are not useful. We presented our approach to various stakeholders in the Netherlands, including hosting providers. The main response was that the metrics were a valid starting point for evaluating hosting security, incentivizing self-regulation and, ultimately, identifying actors for enforcement activities.

The way forward is to improve the methodology. First, we aim to include additional abuse data feeds and more uptime data. We also are improving the identification of hosting providers by using WHOIS data on IP address ownership, rather than AS-level routing data. We are working on techniques to differentiate various hosting provider services, so that abuse rates can take these into account. A further step is to use different aggregation techniques, such as factor analysis. We also aim to further investigate incentives under certain conditions where security metrics can be gamed by providers. Last, but not least, we will undertake more in-depth sensitivity analysis of how the various methodological decisions impact the metric.

It is safe to say that a lot of current claims about hosting providers are based on anecdotal evidence or methods that are not adequately understood. This paper contributes to remediating this shortfall.

## Acknowledgements

This work was funded by NWO under Pr. Nr. CYBSEC.12.003 / 628.001.003, and SIDN ([www.sidn.nl](http://www.sidn.nl)). We would like to thank Paul Vixie and Eric Ziegast for great support, NCSC, the Dutch police, Farsight Security, StopBadware, APWG, Meldpunt, Shadow Server, Abuse.ch, and PhishTank for providing access to their data.

## References

- [1] Brett Stone-Gross et al. “FIRE: Finding Rogue nEtworks”. In: *ACSAC*. 2009, pp. 231–240.
- [2] Davide Canali, Davide Balzarotti, and Aurélien Francillon. “The role of web hosting providers in detecting compromised websites”. In: *WWW*. May 2013, pp. 177–188.
- [3] Craig A. Shue, Andrew J. Kalafut, and Minaxi Gupta. “Abnormally Malicious Autonomous Systems and Their Internet Connectivity”. In: *IEEE/ACM TON* 20.1 (2012), pp. 220–230.
- [4] *HostExploit*. URL: <http://hostexploit.com/>.
- [5] T Moore and R Clayton. “The impact of incentives on notice and take-down”. In: *Managing Information Risk and the Economics of Security* (2009), pp. 199–223.
- [6] C. Wagner et al. “ASMATRA: Ranking ASs providing transit service to malware hosters”. In: *Integrated Network Management*. 2013, pp. 260–268.
- [7] Danny Bradbury. “Testing the defences of bullet-proof hosting companies”. In: *Network Security* 2014.6 (June 2014), pp. 8–12.
- [8] Andreas Pitsillidis et al. “Taster’s choice: a comparative analysis of spam feeds”. In: *IMC*. 2012, p. 427.
- [9] Leigh Metcalf and Jonathan M Spring. *Everything You Wanted to Know About Blacklists But Were Afraid to Ask*. Tech. rep. CERT Network Situational Awareness Group, 2013.
- [10] M Van Eeten, JM Bauer, and H Asghari. “The role of internet service providers in botnet mitigation an empirical analysis based on spam data”. In: *WEIS*. 2010.
- [11] T Moore and R Clayton. “Examining the impact of website take-down on phishing”. In: *eCrime*. 2007.
- [12] Marc Kühner, Christian Rossow, and Thorsten Holz. “Paint It Black: Evaluating the Effectiveness of Malware Blacklists”. In: *RAID*. Vol. 7462. LNCS. Cham, 2012, pp. 1–21.
- [13] Shuang Hao et al. “Understanding the domain registration behavior of spammers”. In: *IMC*. Oct. 2013, pp. 63–76.
- [14] Stjepan Gros, Mislav Stublić, and Leonardo Jelenović. “Determining autonomous systems reputation based on DNS measurements”. In: *MIPRO*. 2012, pp. 1526–1528.