Determining Autonomous Systems Reputation based on DNS measurements

Stjepan Groš, Mislav Stublić and Leonardo Jelenović

***Abstract*—Security on the Internet is a serious problem without satisfactory solution. One problem is at the level of Internet service providers and autonomous systems. This space is highly distributed and without central control, driven primarily by the economic factors. Many solutions have been proposed, with moderate success, concentrating mainly on the Internet routing protocol BGP. We approached this problem with an observation that there is a certain similarity between the Internet's organization at the level of autonomous systems and peer-to-peer networks and thus certain similarity with respect to security issues. In peer-to-peer networks reputation mechanisms are the primary means of protection. We propose similar reputation mechanisms to be applied to autonomous systems. There are many factors that could be used for reputation calculation per autonomous systems, like spam, worms, DoS attacks. In this paper we concentrate only on DNS traffic and propose reputation calculation based on it. Our results show that it is possible to make judgements about entities on the Internet based on the errors found in their traffic.**

***Index Terms*—reputation, security, ISP, autonomous systems, Internet, peer-to-peer networks, DNS**

I. Introduction

It is a fact that today's Internet is plagued with all kinds of security problems. This is evidenced by large quantities of malware, constant attacks, hijacking of network addresses, etc. [1]. The majority of those problems is caused by the Internet's decentralized nature consisting mainly of mutually equal and competitive Autonomous Systems (AS). Within this architecture there is no common oversight body and there are no mechanisms that would punish misbehaving ASes or force them to better control their customers and their own resources in order not to "harm" Internet as the whole. In other words, anything done by any AS has to be justified by an appropriate return on investment (ROI) for the same AS. This explicitly rules out any change done for a common cause as the *common cause* usually doesn't yield direct ROI. Even worse, any network setup and configuration that doesn't affect local customers is acceptable, even if it has negative affects on the rest of the Internet (e.g. [2]). This situation can occur either intentionally, or unintentionally due to the lack of knowledge or carelessness of a person in charge.

In order to solve some of the security problems on the Internet we base our thinking on the premise that the problems are the primary consequence of the decentralized nature of the Internet. This is very similar, though not identical, to a peer-to-peer network. We note that the Internet's architecture at the level of the autonomous systems can be viewed as a form of a peer-to-peer network. An inherent characteristic of the peer-to-peer network is a lack of central authority, just like on the Internet, and this is regarded as a feature, not a bug. Thus, this will not be changed. In a peer-to-peer network similar problems to those on the Internet are being solved using reputation [3]. This analogy suggests that the security measures from peer-to-peer networks might be used by autonomous systems in order to enhance their protection and improve their behavior. This is not the first paper that suggests reputation approach (e.g. [1], [4]), but our view differs significantly as we do not seek to find some absolute values or truths about the Internet as a whole, instead we acknowledge that each entity on Internet see the Internet differently and has a different view about it.

The next premise of our work, and difference from the existing work, is that it would be hard to expect that all, or even majority, of ASes introduce reputation system at once. Thus, there should be some benefit for early adopters. In other words, single autonomous system can measure and monitor other autonomous systems it communicates with and based on those measurements it can determine how good they are. Once significant number of other autonomous systems introduces similar measurements, all those measurements could be combined to form even better reputation value of other ASes.

Once that reputation of ASes is available, ISPs can use this information to better protect themselves and to provide better service to their customers. For example, when there are multiple available AS paths to certain destination ISP can take into account reputation of paths to make better selection. Or, traffic entering ISP can be assigned a different Differentiated Services Code Point (DSCP) [5]. At the moment congestion occurs, traffic can be dropped based on this information. This, in effect, will punish misbehaving ASes. Finally, this information can be made available to different services, like SPAM detection.

To implement this reputation system we hypothesize that the goodness of other ASes could be determined by monitoring what comes from those other ASes to us. Analyzing the incoming traffic (e.g. error packets) the quality of ASes on the path through which this traffic comes can be inferred. In this paper we present preliminary results obtained by monitoring DNS traffic. We also tried to correlate, as best as we could, the obtained results about ASes with the available information on the Internet.

All the authors are with the Faculty of Electrical Engineering and Computing, University of Zagreb, Croatia,
E-mail: {stjepan.gros, leonardo.jelenkovic}@fer.hr, mislav.stublic@gmail.com

Manuscript received MMMMM DD, 2011; revised MMMMMM DD, 2011.

II. Building reputation from DNS traffic

DNS [6] is a well-known and a well-studied protocol. But all the studies done so far tried to gain a global picture about DNS state, and none, that we are aware of, tried to infer anything about the goodness of ASes from which this traffic originates or through which it passes. So, we started by enumerating known DNS errors and analyzed them in order to select only those that we think can tell us something about the quality of the network they come from. Furthermore, we quantified each error by severity. Table I shows the results of that analysis.

TABLE I
DNS errors with estimated severity

Error	Client	Server
Response to erroneous requests	-40	-10
Refused requests	-20	-5
Failed responses (SERVFAIL)	-10	-20
Requests for non-existant domains (NXDOMAIN)	-15	-5
Not implemented RRs	-30	-20
Requests for private addresses (RFC1918)	-50	0
Non-existent TLDs	-60	0
A requests for IP address (A for A)	-60	0
Using port in request 0	-30	-20
Requests with invalid (forbidden) characters	-30	0
Unknown query class	-40	0
Obsoleted and experimental requests	-30	-30

Severity quantifications are based upon the study of what causes them, the amount of occurrences of a particular error found during our own study and also as observed by other authors [7][8]. Also, we note that the client and server side are not equally responsible for a given error so they are not equally punished. For example, when we see *Refused request* then we punish the client (-20) more than a server (-5) since the error is more likely client's fault even though there is a small possibility that it is also on a server side. On the other hand *Using port 0 in request* is more severe when it occurs on a client (value -30) then on a server (value -20).

It should be noted that erroneous packet doesn't mean that the error is on, e.g., server side. For example, when we receive a *Refused requests* packet it could be completely legitimate for a server to send such a response. But we still add a small penalty and count on the scale, i.e. for large number of requests small number of wrongly attributed penalties will go unnoticed.

Reputation itself is calculated in discrete steps of equal and predetermined duration. During each step, n, we collect DNS traffic. Each packet received is classified by AS from which it came from (based on the source address). DNS packets with error in one group (i.e. coming from a single autonomous system) are used to calculate reputation value for the group in the given interval. This value is denoted as R_w. We experimented with several formulas for R_w to study their behavior and to find one which shows the best characteristics. Some of the formulas used include the following representative ones:

$$R_{w1} = \sum r_i \times \beta_i, \tag{1}$$

$$R_{w2} = \frac{\sum r_i \times \beta_i}{A_n}, \tag{2}$$

$$R_{w3} = r_i \times \frac{\sum r_i \times \beta_i}{A_n}, \tag{3}$$

where β_i represents severity quantifications of an error i (as given in the Table I), r_i is the number of occurrences of that error in collected DNS traffic for a given AS in the n^{th} time interval, and A_n is the total number of observed DNS packets (for the same AS and time interval).

Reputation value R_{w1} represents unnormalized reputation (moderated by severity), while R_{w2} represents reputation normalized with the total number of DNS packets in a given time interval. Third reputation value R_{w3} is somewhere between the first two: ASes with more unwanted traffic are more penalized with R_{w3} than with R_{w2} but less than with R_{w1}. In other words, R_{w1} and R_{w3} will point out ASes with lots of DNS traffic with DNS errors, while R_{w2} will highlight ASes with worst ratio of DNS error packets within DNS packets.

In order to determine cumulative reputation in the next time interval, R_{n+1}, that will be used to make judgments about ASes, we were using the following formula:

$$R_{n+1} = \alpha \times R_n + (1 - \alpha) \times R_w, \tag{4}$$

where R_w is one of the R_{w1}, R_{w2} or R_{w3} and parameter α is decay factor, e.g. how fast old reputation is forgotten.

III. Experimental results and discussion

To verify our reputation system and hypotheses from introduction we performed several experiments. Data used in analysis was collected from three points all in different ASes. Some statistics about used data are shown in the Table II.

TABLE II
Basic data about DNS traffic used for experiments

AS	Monitoring period	Packets
AS2108	3-12 Oct 2010	2.53×10^6
AS2108	20-30 Oct 2010	3.68×10^6
AS35549	14-27 Oct 2010	1.12×10^6
AS49788	6-16 Oct 2010	10^5

Within given data we found that appropriate value for α is 0.7 and appropriate time interval is one day. Smaller values have more erratic behaviour, while larger tend to be constant.

Fig. 1 shows cumulative reputations for AS10297 as viewed from AS2108 using all three formulas to determine reputation in current interval, i.e. R_{w1}, R_{w2} and R_{w3}. Since the formulas produce very different numbers, we can only compare them by their behaviour through time, not by absolute values. Initial fall for all formulas is due to initial reputation being set to zero. By the end of the measurement period formulas reach more constant values. Formula R_{w1} has a tipping point at the end of the work week (day 8), and starts rising during the weekend due to lower amount of traffic. As R_{w2} produces reputation values relative to total traffic, the change in total traffic (i.e. during weekends) doesn't change it's behaviour since relative number of errors stays the same.

Using equations (1), (2), and (3) we also determined the worst ASes as seen from each monitoring point. As every equation has it's strengths and weaknesses and tells us different information about AS behaviour all three equations

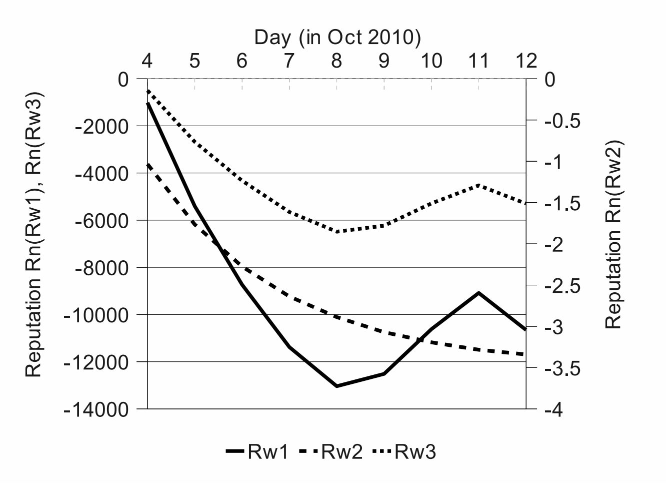

Fig. 1. Reputation of AS10297 as seen from AS2108 calculate from the first data set (3-12.Oct.2010)

are taken into account. Three separate average reputations are calculated for each AS, one for every equation. Average reputations are calculated over a whole measurement period. As we compare average reputations we get three separate worst AS rankings, again, one for every equation. To get worst absolute ASes, all three rankings are taken into account. The final ranking for each AS is created by adding rank as given by every average reputation. For example, if certain AS is 2nd by R_{w1}, 4th by R_{w2}, and 3rd by R_{w3} then it's total score is $2 + 4 + 3 = 7$ and it is ranked by number 7.

Using this method we can find worst ASes for every node where traffic was collected. Table III shows ten worst ASes as seen from different nodes.

TABLE III
WORST 10 ASES BY REPUTATION

AS:	2108	2108	35549	47988
Date:	3-12 Oct	20-30 Oct	14-17 Oct	6-16 Oct
1.	10297	10297	15083	6939
2.	15083	45899	29550	14618
3.	45899	15083	3676	47955
4.	29550	29550	10297	112
5.	14618	22927	3599	15083
6.	3599	8358	5388	10297
7.	5391	14618	7754	3599
8.	17974	17974	17370	22576
9.	6478	45595	8069	4323
10.	9829	4134	21788	29550

AS10297 is the worst AS from the point of view of the AS2108 node. It is Columbus network access point, a company from Columbus, Ohio. Now the AS number belongs to eNET Inc. based in the same town. According to *hostexploit.com*, and *sitevet.com* it is still one of the worst ASes. Among the worst ASes are also AS15083 or Infolink an IT company from Miami, and AS45899 or the VNPT Corp. a telecommunications company from Vietnam.

Data gathering on different nodes was not performed simultaneously on all nodes, although there is a bit of overlapping. However, reputation calculated for two different time periods on the same node within AS2108 suggests that behavior of peer ASes (i.e. ASes from which we receive traffic) is fairly constant, at least over the span of a month. Based on this we can assume similar behavior on other nodes for the same time span. With this assumption we can compare reputations of ASes calculated on different nodes as if they are measured simultaneously. As expected, the reputation picture of the Internet differs based on the measurement point. Still, several of the worst ASes appear throughout all measurement points.

Even though the preliminary results show promise, there is a severe problem of spoofed source IP addresses and attacks on the reputation system itself. This was not taken into account in these measurements and it seems, based on the correlation with available information on the Internet about bad ASes, that it didn't influence results much. Still, we believe that the problem of spoofed addresses can be solved using information available in the BGP. In other words, currently we are punishing apparent source of the erroneous traffic. For future work we left the cases when we are not certain on the source address. In those cases, we could, using BGP, punish path through which the packets had to pass.

IV. CONCLUSION AND FUTURE WORK

In this paper we propose application of reputation system in autonomous systems, as used in peer-to-peer networks, in order to make the actions on the Internet accountable. Presented reputation gathering method is only an example since it is solely based on erroneous DNS traffic. Used methodology for DNS traffic analysis and reputation calculation are promising since ASes with worst reputations are already recognized by other sources. We are aware that the change we propose with reputation system is both very ambitious and demanding in terms of the development resources required. In order to make this more realistic we expect feedback from collogues and, particularly, network operators.

REFERENCES

[1] L. Andersson, E. Davies, and L. Zhang, "Report from the IAB workshop on Unwanted Traffic March 9-10, 2006," RFC 4948 (Informational), Aug. 2007. [Online]. Available: http://www.ietf.org/rfc/rfc4948.txt

[2] R. van den Berg and P. Dibowitz, "Over-zealous security administrators are breaking the internet," in *LISA '02: Proceedings of the 16th USENIX conference on System administration*. Berkeley, CA, USA: USENIX Association, 2002, pp. 213–218.

[3] S. Marti and H. Garcia-Molina, "Taxonomy of trust: Categorizing p2p reputation systems," *Computer Networks*, vol. 50, pp. 472–484, Jul. 2006.

[4] H. Yu, J. Rexford, and E. W. Felten, "A distributed reputation approach to cooperative internet routing protection," in *1st IEEE ICNP Workshop on Secure Network Protocols*, ser. IEEE Computer Society, IEEE Computer Society. Los Alamitos, CA, USA: IEEE Computer Society, Nov. 2005, pp. 73–78.

[5] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An Architecture for Differentiated Service," RFC 2475 (Informational), Dec. 1998, updated by RFC 3260. [Online]. Available: http://www.ietf.org/rfc/rfc2475.txt

[6] P. Mockapetris, "Domain names - implementation and specification," RFC 1035 (Standard), Nov. 1987. [Online]. Available: http://www.ietf.org/rfc/rfc1035.txt

[7] D. Wessels and M. Fomenkov, "Wow, that's a lot of packets," in *PAM 2003: Procedings of Passive and Active Measurement Workshop*, April 2003.

[8] S. Castro, D. Wessels, M. Fomenkov, and K. Claffy, "A day at the root of the internet," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 5, 2008.