Reputation based protection of ISPs

Stjepan Groš
Faculty of Electrical and Computing Engineering
University of Zagreb
Unska bb, 10000 Zagreb, Croatia
stjepan.gros@fer.hr

ABSTRACT

Security on the Internet is a serious and broad problem without satisfactory solution so far. One of the problems faced by the Internet 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. This creates huge problems and threatens the Internet stability and long term growth. Many solutions have been proposed, with moderate success, concentrating mainly on the security of the main Internet's routing protocol BGP. In this paper we note that there is a certain similarity between the Internet's organization and peer-to-peer networks with respect to security. In peer-to-peer systems reputation mechanisms are the primary means of protection. Based on this observation we try to define how reputation could be used on the Internet in order to protect the Internet service providers. The additional benefit is that this system creates economic incentive for the Internet service providers to invest into security even though, by the current measures, this doesn't bring any return of investment.

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: General—*Security and protection*; C.2.5 [**Computer-Communication Networks**]: Local and Wide-Area Networks—*Internet*

General Terms

Security

Keywords

reputation, security, trust, ISP, autonomous systems, Internet, peer-to-peer networks

1. INTRODUCTION

[**?**]

Majority of the security problems on the Internet are caused by it's decentralized nature consisting mainly of mutually competitive and equal Autonomous Systems. Within this architecture there is no common oversight body and there are no mechanisms that would punish misbehaving ASes. Furthermore, the main design principle of the Internet was *dumb network and smart end nodes*, i.e. there is no much state in the network, and overall it is designs avoid to keep state in the core as it complicates design and reduces scalability. The fact is that this principle is the reason for the Internet's great success, and also the source of it's current problems. It's obvious that it would be great to keep this principle as the Internet is further developed and enhanced, but in the same time there is strong requirement to make it more secure than it is now, which implies putting more state into the network. Still, the root cause of the problems that plague the Internet is that it's highly competitive place and almost everything is motivated by the economic gain. Because of that, anything done by an Autonomous System has to be justified by an appropriate return of investment. This explicitly rules out any change done for a common course as the *common course* doesn't yield direct ROI (this is the best example of the *tragedy of the commons* paradox on the Internet). Even worse, anything configured that doesn't affect local customers, but affects the rest of the Internet is acceptable (e.g. [1]). This situation can occur either intentionally, or unintentionally due to the lack of knowledge or carelessness of a responsible person.

The results of such behavior are visible on today's Internet. First, there is a large quantity of unwanted traffic[2]. This traffic, consisting of spam, worms, DoS attacks and similar, is very hard to detect and eliminate, and even harder to trace to it's source in order to stop it. Not only that, but certain Internet Providers even knowingly host generators of the unwanted traffic as this brings them cash. The other seemingly unrelated effect is exemplified by the recent configuration mistake that made YouTube unaccessible[3]. What happened is that one Internet Service Provider made a mistake in their BGP routers configuration and this error propagated via their upstream provider to the global Internet. Thus, their upstream provider also had configuration errors that only now became visible.[1] Related incident also occurred recently where large IP prefix was hijacked[4]. What is common to both, YouTube and the hijack incidents, is that the respective upstream provider didn't block prefixes not owned by the misbehaving provider.

[1]We note that it is an interesting research question whether there were any previous signs that could predict such an outcome?

We believe that the solution for aforementioned problems will not come over the night, they will not be revolutionary and brake everything in order to fix what's broken at the beginning. Neither the solution will be deployed any time soon, and finally, probably there will be a relatively long transitional period. Such opinion is backed by the number of the observations. First, Internet is not anymore technical playground and technical excellence is not anywhere near the first place of priorities. It is used to earn money and thus it's amenable to economic laws and private, selfish, interests! This is somewhat related to the fact that too many people depend on the Internet now and so disruptions are not welcomed as they mean losses! Finally, there is experience of a painful and not yet finished introduction of IPv6 on the Internet after 10 year despite the continuous warnings about IPv4 address exhaustion!

All this has to be taken into the consideration when proposing any solution. Thus, 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 peer-to-peer systems. 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 system. Inherent characteristic of the peer-to-peer system is a lack of central authority, just like in the Internet, and this is regarded as a feature not a bug. Thus, this feature probably wouldn't and couldn't be removed. In a peer-to-peer systems similar problems to those on the Internet are being solved using reputation[5] and trust. This analogy suggests that the security measures from peer-to-peer systems might be used by autonomous systems in order to enhance their protection. This is not the first paper that uses such an approach (e.g. [2, 6]), but our view is broader and encompasses protection of the whole autonomous system and it's infrastructure, not only the routing system, or specific subsystems.

The base premise of this paper is that single autonomous system can measure and monitor all the other autonomous systems it communicates with and based on those measurements in can determine their reputation. Combined with the recommendations obtained from trusted autonomous systems and human operators, all the other autonomous systems can have their *trust* determined. Then, based on the trustworthiness, appropriate actions can be taken with respect to interactions with them. This essentially allows the autonomous system to have better protection with respect to the rest of the Internet and to better schedule it's scarce resources. The additional gain obtained by implementing this protection system is that autonomous systems will have a motive to better run their operations and control their users in order not to earn bad reputation, and thus appropriate repercussions, by other autonomous systems. Finally, we note that by implementing the protection system proposed in this paper we effectively get a *trust-modulated transparency*[7].

The paper is structured as follows. Section 2 reviews basic concepts necessary for understanding the material presented in the subsequent sections. It also presents threat model and finally, enumerates requirements on the presented solution. In Section 3 we describe reputation system and measurement mechanisms that can be used. We also describe protective measures that can be taken based on the calculated reputation value. Then, in Section 4 we propose an architecture of a system built upon the ideas presented in the Section 3 along with the deployment plan for this architecture. The paper finishes with overview of the related work in Section 5 and conclusions and future work in the section 6.

2. PRELIMINARIES

In this section we'll review different concepts necessary for understanding the rest of this paper.

2.1 Autonomous systems

The primary architectural element of the Internet are autonomous systems (AS). They are identified by a *fixed* 16 or 32 bit number. It is important to note that this number can not be easily changed and thus represents fixed identity of any autonomous system. Furthermore, the autonomous systems are described in a databases quired using *Whois* protocol. Autonomous systems are a self contained collections of networks, mostly under single administrative control, running common intra-domain routing protocol. Usually, Internet Service Providers are equal with autonomous systems, though, it doesn't have to be case. For example, when Internet Service Provider delegates parts of it's address space to a customer, the customer should be responsible for the traffic originating from it's part of the address space. Thus, we'll use the term *Internet Service Provider* to mean a network under control of a single entity and responsible by the entity which can be a whole AS or only a part of it, though the major part. Note that enterprises that connect to the Internet but do not provide Internet services fall under this definition as well.

Internet Service Providers exchange control information and data. Control information allows connectivity of autonomous systems and BGP is the main protocol for this purpose. Data on the other hand is generated almost exclusively by the costumers of the ISP.

In this paper we assume that a single ISP is a unit to be protected. The adversaries are other ISPs, their customers and even the customers of the protected ISP itself. The core assets of any ISP are: routers; links between routers; routing protocols, both IGP and EGP; customers and customer's equipment; and any services that are operated by the ISP, i.e. DNS, Web, mail.

Attacks that can be launched against those assets are different denial of service attacks, SPAM, and different attacks on BGP routers and exchange of misconfigured information.

It should be noted that the trust architecture proposed by this paper can also be an attacker's goal. These attacks can be targeted either to disable the reputation service or to influence in some way the computations done by it.

2.2 Reputation Systems

Reputation system is an important part of peer-to-peer networks that helps a node within the network to choose a reliable peer within the same network to transact with[5].

The functionality of a reputation system can be subdivided into three major parts, i.e. information gathering, ranking and taking action.

Information obtained in the gathering phase can be from the direct interactions with the node, or taken from some other node that had direct or indirect experiences with target node. When the node uses information collected from some other node, we are talking about *receiving recommendations* from the other node. In the ranking process, the node tries to rank all the other nodes according to expected reliability and, finally, based on the ranking process it selects peers to transact with and takes an action.

In peer to peer systems adversaries powers include *traitors*, *collusion*, *front peers*, *whitewashers* and *denial of service*. Traitors behave properly for a period of time to build good reputation and then use this reputation to misbehave. Collusion is grouping of misbehaving peer that act as a group. Front peers help other peers to gain good reputation and then those other peers perform attacks. Whitewashers are peers that change identity in order to get rid of negative reputation. Finally, DoS attacks are not specific to peer to peer networks, but they can also be targeted against reputation system itself.

2.3 Requirements

The success of a new technology depends on it's deployment. The solutions that require all players to implement given technology in order to be useful are predestined to failure. Similarly, if deploying some solution doesn't bear direct benefit to the one using it then the solution will fail also. Thus, the main requirement we set for the solution is that (*i*) it's useful if only locally deployed, and (*ii*) it brings direct and observable benefit to a user!

Next, we note that there are attacks on the Internet that are effective within minutes. The solution proposed in this paper doesn't counter them directly, but rather, indirectly. In other words, the goal is to prevent attacks, before they have a chance of exploding. At a start of the deployment these attacks would still be possible, but as the user base grows we believe that it will become ever more harder to initiate new attack.

Furthermore, the use of reputation systems in the peer to peer networks is much more advanced than our proposal here. For now we'll only concentrate on cases where single ISP is using this mechanism for the protection. The more complex cases enabled by better deployment are left for the later phase of this project. This is justified by the expectation that first deployments will be sporadic and organic.

3. BUILDING THE REPUTATION OF ISPS

In order to build a reputation of all the other ISPs, the reputation system performs it's job in three phases, i.e. information gathering, reputation scoring and ranking, and taking action. We first describe the method used to calculate reputation values based on the current reputation and observed behavior of the target ISP. Then we give an overview of few measurements we analyzed thus far. Finally, the protection mechanisms the ISP has on it's disposal are reviewed.

As we already discussed, the identity system forms a basis of a reputation system since the reputations are bound to identities. The identity system is based on the fact that each ISP is uniquely identified by it's autonomous system number. Furthermore, there are public databases enumerating assigned networks to autonomous systems. Thus, we'll use autonomous system numbers for identifiers. There is potential problem with this approach that warrants further research. Namely, certain ISPs delegate parts of their IP address space to customers and publish those delegations in appropriate databases (e.g. whois). When such customers are a source of the problem it might not be good to blame ISP.

Possible problem for the proposed identity system could be DoS attacks that use spoofed IP addresses. Care should be taken in order not to blame innocent ISP simply because we blindly believed to source address. Thus, sources have to be verified in some way. Still, in case of spoofed sources operator of the attacked ISP can try to determine actual source using some form of IP traceback mechanisms[8]. In case that some autonomous system between destination and real source does not support traceback, or blocks it for whatever reason, then it will be credited for the attack and consequently it's reputation will be degraded. This, presumably, will motivate ISPs to cooperate in such situations.

3.1 Reputation scoring and ranking

Based on the collected data reputation of each autonomous system should be computed. First we have to define what value the reputation system will have, and also we have to define an initial value for a new entity. There are three approaches to this problem. The first one is with low or no reputation at all. The pro for this variant is that today AS can be entity with enough money and connectivity and there are no special requirements to enter ISP space. Special requirement could be adequate number of trained people, or something similar. On the other hand, new ISP hasn't done anything wrong yet and thus, it's not right to place it into the same reputation level as all the ISPs that continuously harm the Internet. Thus, we could start with maximum reputation value and lower it for each mistake done by the target ISP. Still, this approach is not good also. Suppose that we want to reward an ISP that implemented some optional feature, like S-BGP. In that case, if the given ISP already has maximum reputation value as it's behaving very good, there is no way to implement reward process. Thus, we'll use third approach where each new ISP has middle, or average, reputation. For easier computations we'll select that the reputation is in the range of $[-1, 1]$, then new, or previously unseen, ISP will have reputation value of 0.

The next decision is related to exact calculation of the reputation value. The general principles that should be obeyed in the calculation are:

1. The past has to be taken into account when calculating a current value for the reputation.

2. When there is neither good nor bad information the reputation should, with time, asymptotically approach default value.

3. Different measurements should be kept separate, e.g. measurements of SPAM originated from an AS should be separate from measurements of the BGP stability to an ISP.

Based on those principles we propose the function similar to the smoothed average function used in the TCP. In other words, for each new sample of some parameter we calculate the new value as follows:

$$r_{i+1} = \alpha \times r_i + (1 - \alpha) \times m_i \quad (1)$$

In the formula α determines how fast the past is forgotten, or equivalently, how much influence the new values have. The right value will be determined based on the measurements we plan to conduct. The value m_i is a snapshot of a value of some parameter. This value should be further multiplied by some constant factor that reflects trustworthiness of the measurement device.

To accommodate the requirement that the reputation value, in case when there is neither good nor bad information, asymptotically approaches neutral value, we propose that the new value of the reputation is recalculated when either there is a new measurement, or, some predefined time interval passes. In later case, the same formula is used, but the sample is set to zero.

Finally, the third requirement states that we should keep all the parameters that influence reputation separately and combine them on demand. For this reason each observed parameter is kept separately and for each ISP we are aware of. The given equation will most probably be vectorized as there are multiple measurements, and thus, reputation will be a vector. To make some decisions based on reputation, two characteristic values could be taken, average of all the components and standard deviation. Of course, there are plenty of possibilities and this requires further research.

3.2 Information gathering

In the information gathering phase we are collecting as much information about peers as possible. Usually, this is done via direct or indirect interactions, but at the beginning of the deployment of the mechanisms proposed in this solution, direct interactions will dominate. We expect that as the method becomes more popular there will be possibility to also use experiences that others have with certain peers. In that case simple extension in a manner of [9] can be used.

There has been multitude of different measurements on the Internet regularly published on different conferences. Good number of those measurements is used as an isolated indicator of Internet's health and ISPs quality. We plan on integrating as much of those measurements as possible in order to gain insight as accurate as possible into each ISP behavior and operational procedures. Because of the size constraints of this paper and relative beginnings of the project we'll only present several examples of measurements. For each one we'll describe proposed way of including it as an input into reputation calculation function.

Potential sources for information gathering can be grouped into the following categories:

- Probes sent by the Internet Service Provider
- Monitoring of all the other Internet service providers
- Operator intervention

3.2.1 SPAM

Spam values are recalculated at periodic time intervals. At the start of each interval all the ASes have counters initialized to zero. During the interval period votes from SPAM detection sources are collected. Sources are mail servers, either those belonging to the ISP itself, belonging to the ISP's customer or someone that ISP believes to. Each mail server for each SPAM reports source IP address of the mail server that originated SPAM. This address is converted to AS and the counter for the AS is incremented by the trustworthiness of the reporting server (the value in the range $(0, 1]$) and added to an existing value. At the moment the timeout expires, for each AS new SPAM reputation is determined using the formula similar to 1:

$$s_{j,i+1} = \alpha_{SPAM} \times s_{j,i} - (1 - \alpha_{SPAM}) \times \frac{1}{1 + e^{-c_j}} \quad (2)$$

The index j iterates over all networks known to the ISP, the $s_{j,i+1}$ is a new SPAM penalty value for the j AS, α_{SPAM} is a forgiveness factor, $s_{j,i}$ is the previous value and c_j is the counter's value for the given AS.

3.2.2 BGP measurements

Several BGP parameters could be used to determine the reputation of ISP. The first one is stability of it's routes, i.e. the number of ANNOUNCE/WITHDRAW messages. We assume that frequent changes of routes signal error or misconfiguration. There are already sites that measure number of updates per AS, and number of updates per prefix per AS[**?**].

The problem with the BGP is that sometimes it's not possible to determine who exactly caused the error. For example, if the router receives withdrawal of some network that is reached over several autonomous systems, than each autonomous system could be guilty for the error. This warrants further research.

3.2.3 Operator's intervention

There are important experiences when dealing with other ISPs that can not be automatically measured. Examples of such experiences are successful or unsuccessful mail exchanges to well-known contact mail addresses[10]. Furthermore, even knowing operators in other ISPs might influence the reputation of those ISPs. Finally, there are possibilities of transient errors by reputation system. In such cases it is envisioned that operators have several possibilities to influence reputations. First, they can completely change current reputation value of certain ISP. They can also cast vote, negative or positive, about some ISP that is added to current

reputation just as any other measured value. This vote can even be pondered based on the privilege level of the operator himself.

3.3 Actions based on reputation

Upon determining trust level of different autonomous systems on the Internet, the Internet Service Provider can use this information in order to protect itself and it's customers. In this section we review some protective measures we envision that will be available to the ISP and discuss potential problems that can arise.

The basic mechanism that the ISP has on it's disposal is treatment of traffic that flows through it's routers. The idea is that the less trustworthy some autonomous system is it is more likely that it will originate malicious traffic. Thus, by default, traffic that originates from more trustworthy sources should receive better treatment that the traffic from less trustworthy sources.

In order to implement this policy the whole trust range is subdivided into set of classes, e.g. 8 classes. The highest class is reserved for the traffic that originates from the ISP itself or from it's customers. All the other classes are for the rest of the Internet and anything received at the boundary of the ISP is checked against trust level of the originator. Based on the result of those checks each packet is marked with the appropriate DSCP value[11]. Based on such markings in the traffic all the routers have information about trustworthiness of the source embedded into the packet itself which doesn't slow down forwarding process. Not only that the routers have different behavior based on the DS markings, but so could also firewalls and intrusion detection systems. Firewalls can have more strict rules for less trustworthy traffic, while intrusion detection systems can have larger rule databases and/or lower thresholds for less trustworthy traffic.

There are two potential problems: (i) possible traffic starvation, and (ii) scalability. In case that two sources of different level of trustworthiness have constant traffic and compete for the single link, it could happen that the less trusted source is starved. This should be prevented meaning that appropriate queueing management and service functions have to be used in the routers. On the other hand there is potential for scalability problems at the ingress where traffic marking should be done. Namely, for each packet source has to be determined in order to determine trust level and consequently traffic class. As there is a large number of source networks this is a potential bottleneck. We propose two different behaviors. For ISPs with large quantities of data the alternative method would be to mark traffic based on the ingress link. In that case the aggregate of all the sources on the given link should be used.

Another set of actions can be performed by directly querying trust server for reputation of different sources. There are the following possibilities:

- Anti SPAM software can adjust threshold level based on the trustworthiness of the mail's source. This can go to the extreme of not accepting mail unless it's coming from the mail server that is in an ISP of some minimum trust level.
- Servers can base their decisions to answer queries, or even to allocate resources, based on the trustworthiness of the client's source.
- When BGP routers calculate new paths, they can take into the account trustworthiness of the path to the destination. If, for example, there are two paths to the destination D, one going through 3 autonomous systems, and the other one going through 5 autonomous systems, the longer one can be selected based on the fact that those 5 ASes are more trustworthy that the other 3.
- Thresholds in the BGP routers to prevent route flapping can be adjusted based on the trustworthiness of either the source, or the paths that are causing the flapping.

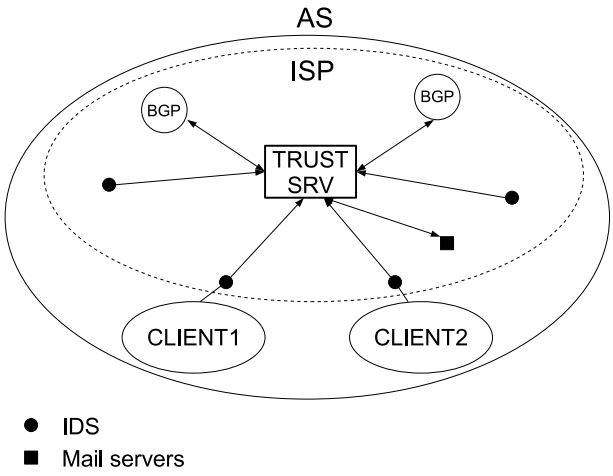

Figure 1: Major components in the trust zone

The majority of the aforementioned actions requires changes in the end nodes software which makes it more intrusive and thus harder to introduce.

4. SYSTEM IMPLEMENTATION

The architecture of a planned reputation system is shown in the Figure 1. The central component of the architecture is the *trust server*. In each trust zone there is at least one trust server, but, for the security, scalability and reliability purposes there could be any number of trust servers distributed on strategic points throughout the zone. *Trust zone* contains all the networks that share identical recommendation system, view of other autonomous systems, and the same mappings of source trustworthiness into traffic DSCP codepoints.

There are several tasks performed by the trust server. First, it enumerates objects for which it maintains trust levels. Then, it collects evidence of positive and negative behavior of the objects. Based on the collected data it calculates

reputation and trust of the objects. Finally, it disseminates trust information to interested parties.

For a purpose of determining trust level of objects monitored by trust server there are measurement nodes spread through the trust zone, but also there could be measurement nodes in another zones of the interest. This possibility we do not discuss at the moment. The purpose of the measurement nodes is to obtain relevant parameters for the trust server. The functionality of measurement nodes can be performed by specialized nodes, or the hosts and routers can have appropriate support and send updates to trust server. A single measurement node can send measured parameters to multiple trust servers.

In the Figure 1, trust server communicates with the BGP routers in order to collect relevant information from them and to adjust traffic classification for the data entering service provider infrastructure. It also relies on Intrusion Detection Nodes that serve as a form of Network Telescope[] which captures data from other autonomous systems. It should be noted that there are also IDS nodes that monitor customer traffic. Finally, there is connection between mail server and the trust server where mail server provides information about spam sources, while trust server provides trustworthiness data to mail server.

Trust servers communicate using *recommendation protocol*. Servers querying trust server for information about specific sources (e.g. mail server querying trustworthiness of some source that just sent an email) also use the recommendation protocol. It is envisioned that this protocol will be developed by the IETF.

5. RELATED WORK

There are plenty of measurements done to characterize the behavior of the Internet, it's components and users. All of them could be treated as related work. The main difference is that those measurements are not planned to be consolidated! Few proposals are based on trust and reputation for the protection of autonomous systems. We review some of them here.

The PATRICIA architecture[9] is based on the premise that all the edge networks could cooperate in defending the Internet. There are few differences with respect to our proposed architecture. First, PATRICIA is designed only for edge networks, while our architecture doesn't depend on a placement of an autonomous system. Secondly, PATRICIA assumes implicit trust between different edge networks, and reliance on action by source network. As we already discussed in the requirements section, we can not rely on others to do something without economic incentive, which, in this case, is not present.

The second similar proposal is described in [6]. There are also few key differences compared to our approach. First, they propose building trust-based peer-to-peer overlay network. Secondly, they only take into account adjacencies of autonomous systems. Finally, though they claim that their proposal is incrementally deployable, but it's hardly useful for a single ISP.

6. CONCLUSION AND FUTURE WORK

In this paper we propose application of reputation system, as used in peer-to-peer networks, in order to make the actions on the Internet accountable. Furthermore, we envision that measurements regularity published in scientific papers will be integrated in order to assess the security state in the Internet. This is not an easy task and this paper is a first step to either accept or reject this idea. We are aware that the change we propose 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.

Our first step is to build a system that will implement the functionality of trust server. This is necessary as it will allow us to experiment with different scoring systems and reputation calculations.

7. REFERENCES

[1] 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.

[2] 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

[3] "YouTube IP Hijacking," Feb 2008. [Online]. Available: http: //www.merit.edu/mail.archives/nanog/msg06299.html

[4] "Latest instalment of the "hijacked /16s" story," Jun 2008. [Online]. Available: http://mailman.nanog.org/pipermail/nanog/2008-June/001416.html

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

[6] 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.

[7] D. Clark, K. Sollins, J. Wroclawski, D. Katabi, J. Kulik, X. Yang, R. Braden, T. Faber, A. Falk, V. Pingali, M. Handley, and N. Chiappa, "New arch: Future generation internet architecture," http: //www.isi.edu/newarch/iDOCS/final.finalreport.pdf, Dec. 2003.

[8] A. Belenky and N. Ansari, "On IP Traceback," *IEEE Communications Magazine*, vol. 41, no. 7, pp. 142–153, 2003.

[9] L. Wang, Q. Wu, and D. D. Luong, "Engaging edge networks in preventing and mitigating undesirable network traffic," in *3rd IEEE Workshop on Secure Network Protocols*. IEEE, 2007, pp. 1–6.

[10] D. Crocker, "Mailbox Names for Common Services, Roles and Functions," RFC 2142 (Proposed Standard), May 1997. [Online]. Available: http://www.ietf.org/rfc/rfc2142.txt

[11] K. Nichols, S. Blake, F. Baker, and D. Black, "Definition of the Differentiated Services Field (DS

Field) in the IPv4 and IPv6 Headers," RFC 2474 (Proposed Standard), Dec. 1998, updated by RFCs 3168, 3260. [Online]. Available: http://www.ietf.org/rfc/rfc2474.txt