Complementing E-Mails with Distinct, Geographic Location Information in Packet-switched IP Networks

Stephan Kubisch[1], Harald Widiger[1], Peter Danielis[1], Jens Schulz[1],
Dirk Timmermann[1], Thomas Bahls[2], Daniel Duchow[2]

[1] University of Rostock
Institute of Applied Microelectronics and Computer Engineering
18051 Rostock, Germany
Tel./Fax: +49 (381) 498-7276 / -1187251
E-mail: {stephan.kubisch;harald.widiger}@uni-rostock.de
Web: http://www.imd.uni-rostock.de/networking
[2] Nokia Siemens Networks GmbH & Co. KG
Broadband Access Division
17489 Greifswald, Germany
Tel./Fax: +49 (3834) 555-642 / -602
E-mail: {thomas.bahls;daniel.duchow}@nsn.com
Web: http://www.nokiasiemensnetworks.com

Abstract. Although the Internet has developed into a mass-medium for communication and information exchange over the last couple of years, many problems still exist regarding security and anonymity. One of these Achilles' Heels is *spam*. Electronic mail (e-mail) has become one of the most used communication mechanism. It is absolutely easy to use and cost-effective. Unfortunately, the simplicity and effectiveness of e-mail are also major drawbacks. Without additional effort, Internet users' mailboxes are flooded with unsolicited, bulk e-mails of many different flavors; mostly without any chance to identify the true origin. Thus, most anti-spam techniques rely on spam detection using large filter data bases and pattern matching functions but cannot identify the trustability of the sender. To overcome this lack of security and trustability, a new concept—*Trust-by-Wire*—is introduced as well as a mechanism called *IPclip*, which provides the basic means for enhanced e-mail security. The main idea is to guarantee Trust-by-Wire in packet-switched IP networks by providing trustworthy location information along with every IP packet. The availability of suchlike information in conjunction with some trust management is discussed in the light of a conceptual framework, which allows for reliably tracing the geographic origin of e-mails. IPclip furthermore provides an additional trigger for existing spam classification and filtering systems. Thereby, IPclip can neither be fiddled by benign e-mail users, nor by trojaned bots, nor by spammers.

Keywords. E-mail, Simple Mail Transfer Protocol, Internet Protocol, Tracing Spam, Location Information, Trust-by-Wire.

List of Abbreviations

AN	Access Node, page 9
BRAS	Broadband Access Server, page 9
CLIP	Calling Line Identification Presentation, page 7
CPE	Customer Premises Equipment, page 9
DDoS	Distributed Denial-of-Service, page 3
DKIM	DomainKeys Identified Mail, page 5
DNS	Domain Name System, page 5
DoS	Denial-of-Service, page 3
DSL	Digital Subscriber Line, page 9
DSLAM	DSL Access Multiplexer, page 9
FPGA	Field Programmable Gate Array, page 22
Gbit/s	Gigabit per Second, page 22
GPS	Global Positioning System, page 7
HTML	Hypertext Markup Language, page 5
ID	Identifier, page 5
IP(v4/v6)	Internet Protocol (Version4/6), page 3
IPclip	Internet Protocol Calling Line Identification Presentation, page 7
ISDN	Integrated Services Digital Network, page 7
ISP	Internet Service Provider, page 5
LI	Location Information, page 7
MTA	Mail Transfer Agent, page 5
MUA	Mail User Agent, page 6
MX	Mail Exchange, page 13
PF	Peering Flag, page 8
PSTN	Public Switched Telephone Network, page 3
RF	Removal Flag, page 8
RFC	Request for Comments, page 14
SCA	Subscriber Catchment Area, page 10
SF	Source Flag, page 8
SMTP	Simple Mail Transfer Protocol, page 3
Spam	SPiced hAM, http://www.hormelfoods.com/brands/spam, page 3
SPF	Sender Policy Framework, page 5
Spit	Spam over Internet Telephony, page 22
TBW	Trust-by-Wire, page 4
TCP	Transport Control Protocol, page 6
TF	Trustability Flag, page 8
TLV	Type-Length-Value, page 8
UBE	Unsolicited, Bulk E-Mail, page 4
UCE	Unsolicited, Commercial E-Mail, page 4
VoIP	Voice-over-IP, page 22
WHOIS	“Who is?” Domain Name Lookup Service, page 20

1 Introduction

What do quota warnings, false positives, and annoying advertisement have in common? They all are likely to be the result of *spam*, which is one of the modern Internet's Achilles' Heels. The original Internet has grown from a pure scientific network into a full-blown world-wide information and communication medium [1]. But the Internet's open structure is one reason for its frequent misuse. Unfortunately, a "dark side" of the Internet emerged. This has, among others, the following reasons:

- The Internet's complexity and therewith anonymity are increasing. Black sheep can hide easily. Decades ago, the Internet has been a trustable environment. But it has developed into a mass-medium.
- Aged protocols, which have not been designed for such a large community [2] in the first place, show deficiencies. The Simple Mail Transfer Protocol (SMTP) for example, which lacks client authentication, has been used without any concerns. Nobody could foresee the spam problem. Today, the drawbacks and backdoors of the early protocols [3, 4, 5] are increasingly exploited and misused.
- Packet-switched networks do lack Trust-by-Wire. In circuit-switched networks, e.g., the Public Switched Telephone Network (PSTN), a fixed line directly references the caller, whereas this direct interrelationship is not given in todays packet-switched networks. Current and future access networks and the Internet *are* flavors of packet-switched IP networks.

Due to the success of the Internet, e-mail is one of the most used digital communication media today. Millions of people can be reached in a fraction of the time needed for written letters as well as with a fraction of the costs. The first spam was generated in 1978 by Gary Thuerk sending e-mail advertisements to 600 recipients [6]. Since then, e-mail has been abused by so-called spammers not only for advertisement but for Denial-of-Service (DoS) attacks and the propagation of malware or forbidden content as well. By now, spam traffic represents a huge part of the Internet's overall e-mail traffic and supersedes the amount of good e-mails by a huge factor [7, 8]. Approximately, 90% of all e-mails can be classified as spam.

Furthermore, problems and troubles like spam also lead to a reduced belief and confidence of the users in e-mail services and the Internet as central communication medium [9]. However, since e-mail *is* one of the most wide-spread communication media—even cellular phones are capable of receiving and sending e-mails—most e-mail users have learned to live with the spam problem by using different types of filtering technologies or by just ignoring it. Although a broad range of mechanisms for spam detection and prevention exists, there is no 100%-solution out there. Most anti-spam measures are only of reactive nature. Mostly, some kind of filter technology in combination with up-to-date blacklists is used. But even these mechanisms are prone to attacks as the latest headlines show (google for, e.g., *"Spamhaus Under Another Distributed Denial-of-Service attack (DDoS)?"*).

This paper presents the Trust-by-Wire (TBW) concept and the IPclip mechanism in the light of tracing and detecting e-mail spam. A conceptual framework is sketched and discussed regarding its feasibility and usefulness. TBW and IPclip are thereby not intended to replace but to complement existing anti-spam mechanisms. On the one hand, e-mails can reliably be traced with reasonable accuracy to identify the sending host. On the other hand, the provided information is a supplementary trigger for the classification of incoming e-mails, e.g., using tools like SpamAssassin [10].

The remainder of this paper is organized as follows: Section 2 shortly describes the nature of spam and spammers. In the overview in Section 3, we briefly revisit the state-of-the-art in anti-spam mechanisms. Section 4 introduces the general TBW concept and the IPclip mechanism. The use of TBW and IPclip in an anti-spam use case is discussed in Section 5. Some aspects of future work are highlighted in Section 6 before concluding the paper in Section 7.

2 Spam & Spammers[1]

During the last years, many different kinds of e-mail spam have emerged. The latest developments of spam as well as the percental distribution of different spam types are analyzed in [8]. In [11], an extensive analysis on e-mail traffic patterns is presented.

Unsolicited Bulk E-Mail: The most prevalent type of spam is called UBE—unsolicited *and* bulk e-mail. It is sent unrequested to a huge number of recipients. Various subtypes have emerged. Scam e-mails typically contain cheating or deceit attempts. The objectives of phishing e-mails are identity theft and obtaining private user data like logins & passwords or credit card numbers by fraud. Hoax spam tries to communicate false reports and canards. Chain letters with charitable background belong to hoax as well. An unsolicited commercial e-mail (UCE) is nothing more than advertisement to promote dubious, “special” offers and cheap products, drugs, or forged branded articles.

Collateral Spam: Another type of e-mail spam is collateral or secondary spam. Suchlike spam e-mails occur either as response to an incoming spam e-mail or due to malware. The responses are sent back to noninvolved, third persons due to manipulated return addresses. Malware could have infected a host, which now autonomously sends UBEs using that host’s internet connection and IP address. Furthermore, suchlike trojaned hosts may scan the local directories for address books and valid e-mail addresses.

In most of the cases, spammers do only follow financial goals. To be successful, a spammer has to send millions of e-mails in a very short time to many different people. Reasons for that are:

- Anti-spam mechanisms, as summarized in the next section, already filter a considerable part of the spam e-mails.

[1] Readers familiar with these topics might skip Section 2 and 3.

- Many of the misused e-mail addresses do not represent active and valid e-mail accounts (any more).
- Only a minimal part of the recipients does react on the ambiguous offers. But it is still enough to keep spamming profitable.

Thus, spammers are technically well skilled but also impatient—and spammers lie. Different techniques are used to obscure and manipulate the true origin of the e-mail or the e-mail traffic. In [12], Boneh details examples like manipulated sender information in the e-mail header or the exploitation of open proxies and mail relays. Typical routes that ham and spam usually take are sketched in [13]. While direct spamming and open relays & proxies have widely been used in the early days, networks of trojaned hosts (bots, zombies) of up to multiple 100.000 bots are responsible for a large part of spam today [14]. Thereby, bot networks are not only used to generate spam but to drive DDoS attacks as well.

Unfortunately, it is often not known to the mainstream user that sender information can easily be forged and replaced with arbitrary entries. To let the e-mail's look appear trustworthy to the victims and to impair anti-spam filtering tools, spammers exploit HTML features, apply forged images, or jumbled letters.

The consequences of spam are not only high monetary losses. Anti-spam tools must be maintained to be up-to-date and the sorting and reading of spam e-mail costs a considerable part of working time. Furthermore, Internet Service Providers (ISPs) still offer pay-scales, which account for every transmitted byte. Thus, benign users pay for every byte of spam they send & receive.

3 Anti-Spam Mechanisms

Principally, four different classes of anti-spam measures can be found today. (This overview is far from being comprehensive.)

Sender authentication and reputation: Wong gives a comprehensive overview on sender authentication and reputation mechanisms in [15]. Common authentication approaches are the Sender Policy Framework (SPF) Project [16] and DomainKeys Identified Mail (DKIM) [17]. The SPF Project actually consists of two slightly different parts—the literal SPF [18] and Sender-ID [19]. Both use special DNS records to validate sender information. SPF authenticates sender information of the e-mail envelope. Sender-ID, which has its roots in Microsoft's Caller-ID for E-Mail, authenticates sender information of the e-mail header. By contrast, DKIM authenticates the content and various header fields of an e-mail using asymmetric encryption. The recipient or the receiving mail transfer agent (MTA) uses a public key to check the e-mail's credibility.

Black lists and white lists are mostly used to check the reputation of a domain or an IP address[2]. While black lists contain IP addresses of known spammers, which are used to block the respective incoming e-mails, white lists contain address information of benign hosts, which should not be blocked. Mostly, they

[2] A comprehensive collection of B&W lists can be found at http://spamlinks.net.

are maintained by non-profit organizations [20]. Their advantages are fast, near-realtime lookups and good up-to-dateness. Unfortunately, they might also contain false positives & negatives and cannot guarantee 100% protection. However, recent and increasingly frequent attacks against anti-spam lists to compromise their availability prove the importance of B&W lists.

Content analysis: Analyzing the content of header and body of an e-mail is another main pillar in the fight against spam. Although spammers are getting more and more sophisticated in forging the appearance of e-mails, filter technologies still detect a considerable amount of spam. A typical filtering tool, which is implemented on MTAs, is SpamAssassin [10]. Most mail user agent (MUA) software like Mozilla Thunderbird or Microsoft Outlook feature adaptive junk filters to classify incoming e-mails. In [21], a profound review of filter technologies is given.

Throttling e-mail delivery: Another flavor of anti-spam measures is to throttle the e-mail delivery process. Thereby, the forwarding of an e-mail is delayed in the one or the other way. Mostly, these extra seconds are nonrelevant for mainstream users. But they hit a tender spot of spammers, which are impatient by nature. Nevertheless, there are also some types of ham, which may suffer from the additional delay, e.g., when pressing e-mails must be rolled out or in case of emergency. Typical examples for delaying the delivery process are graylisting [22] and tarpits (also know as teergrubing) [23].

Active measures: This type of anti-spam measures is mostly used by experts or expert hobbyists. Scam baiting describes the process of deceiving scammers themselves by answering to spam e-mails, pretending interest, and finally unmasking and exposing the scammer [24, 25]. A promising approach to catch, monitor, and analyze malware are so-called honeynets [26]. Prepared, purpose-built networks and hosts attract all flavors of e-mails including spam—besides a multitude of other threats and malware. Using different analysis tools, knowledge on the behavior and characteristics of spam and malware is collected.

All the techniques mentioned above mostly relate to the use of SMTP, since it is by far the most used e-mail protocol. Some alternatives exist for e-mailing but they are not common, for example Message Submission via TCP-port 587 [27, 28], which closes some loopholes in the SMTP protocol. But ISPs do not dare to block port 25 for all their customers on a network-wide scale.

Furthermore, every e-mail user should hide his own e-mail addresses and address books [29]. Address information should not be published on a website where it is readable for harvester bots that automatically filter website content for e-mail address syntax. E-mails should be published in some encrypted format only. Subscription to mailing lists, e.g., for newsletters, should be done with care as well. So-called confirmed opt-in techniques are to be preferred.

The first two types of anti-spam mechanisms have achieved wide acceptance and are mostly used in some combination. A general overview about anti-spam measures is given in [30]. A typical serial connection of various anti-spam techniques is detailed in [31].

4 Trust-by-Wire & IPclip in General

The TBW framework and the IPclip mechanism have been developed totally decoupled from potential use cases like the one addressed in this paper. Thus, this section provides a brief overview about the general IPclip mechanism before adapting it to the spam use case in Section 5. We refer to [32] for detailed information on the general IPclip mechanism. Another use case addressing protection from phishing attacks is presented in [33].

The name *IPclip* is derived from the CLIP functionality (Calling Line Identification Presentation) in ISDN (Integrated Services Digital Network) telephone networks. CLIP is an optional feature to submit the calling number to the telephonee to present it on, e.g., a display. This way, the callee can identify the caller. In case of packet-switched IP networks, the IP address of a user cannot be treated as equivalent to a fixed line telephone number. The reason is, as already mentioned in the introduction, that an IP address does not necessarily identify a distinct physical line. Furthermore, IP addresses *do not* allow any conclusions on the geographic location of a packet's origin. In contrast, fixed line telephone numbers *do* have a well-defined and known origin. The original idea and the name of the CLIP feature in classical ISDN telephone networks are thus adapted in our TBW framework for packet-switched IP networks. With *Trust-by-Wire*, we describe a direct interrelationship between some flavor of user-ID, e.g., a network address, login name, or phone number, and the physical line or geographic location of that user. In other words, TBW relates to unambiguousness and trustworthiness in packet-switched IP networks. From a technical perspective, IPclip is a completely novel mechanism and cannot be compared with the classic ISDN CLIP.

4.1 Why the Internet Protocol and what Kind of Information?

An Internet user and his actual geographic position can be identified with IPclip using a tuple of information consisting of the customer's current IP address and some additional information. As IP addresses do not clearly reference to the users' locations, reliable location information must be included in the additional information. Preferably, standardized data formats should be used for it in order to ensure global interoperability, which is essential in the Internet. Due to its global availability, the GPS (Global Positioning System) data format is used to encode geographic location information [34] at the moment. The sum of all additional information—in the following just specified as location information (LI)—is used for analysis, classification, or stimulation of further actions.

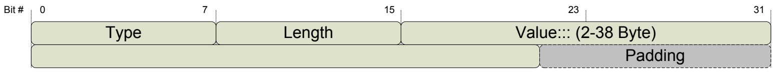

Fig. 1. TLV-structure of an IP option

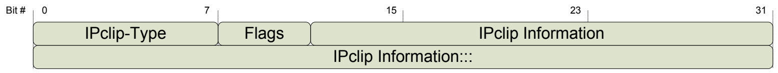

Fig. 2. Structure of an IPclip option inside the value field of an IP option

To provide LI on a global scale, an optional data field is inserted into every IP packet. The reason is that IP is the central protocol in the Internet. IP provides end-to-end connectivity between users, service providers, and network nodes in general. Besides, structure and size of optional fields inside IP, so-called IP options, are standardized [35]. This way, the IPclip mechanism is a standard-compliant solution for the delivery of supplementary LI. Every IP-capable device can either analyze and processes IP options or ignore them. But in any case, devices must at least be able to parse and skip IP options for reasons of interoperability. Next to the feature of adding additional LI into packets, the whole mechanism can be configured to remove suchlike IP options. This may be necessary if Internet users do not want to receive or are not allowed to receive sensitive information about the geographic origin of IP traffic. In these cases, the use of IPclip is totally transparent. However, this depends on the application and regional policies.

The new IP option shows the typical TLV structure (Type-Length-Value) as sketched in Figure 1. The TLV structure must be understood by every IP-compliant network device. The type field is divided into a 3-bit field for various flags and a 5-bit IP option number. For prototyping, we have chosen 26 as option number for IPclip as it is not in use otherwise [36]. Length denotes the IP option length including type and length field. The value field of the new IP option contains the IPclip option. Figure 2 shows the structure of an IPclip option. The IPclip type field denotes the kind of information this IPclip option contains, e.g., GPS coordinates. The option information field contains the actual geographic information, which depends on the IPclip type. The 4-bit status field contains four flags:

- Two flags, the Source Flag (SF) and the Trustability Flag (TF), are used for trust management (see Section 4.3 and Table 1).
- The third flag, the Peering Flag (PF), is used to indicate that an IP packet and IPclip LI arrived from a different ISP's domain.
- The fourth flag, the Removal Flag (RF), is used to indicate that the IPclip LI can be removed on the egress AN before delivering the IP packet to the destination host. It is set on the ingress AN of the source host.

The addition of LI including its analysis and verification raises different important questions:

- Which is the place within the network infrastructure where the LI to be added is available?
- Which is the place within the network infrastructure where this LI can be added into IP packets?

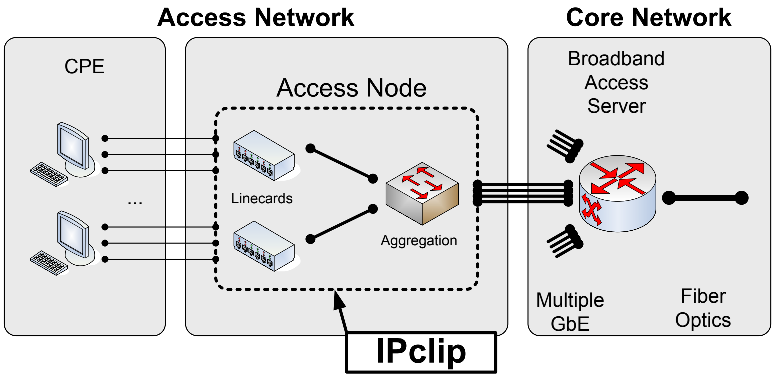

Fig. 3. Network structure with CPEs, access network, and core network.

- How can a trust relationship and a certain degree of credibility be described and how can it be ensured when analyzing and validating the additional information?

4.2 IPclip's Position within the Network Infrastructure

Network ingress—also known as access network—is the most reasonable place where LI can be added and verified. Access networks comprise Customer Premises Equipment (CPEs) as well as so-called access nodes like IP DSL Access Multiplexers (IP DSLAMs) [37]. Usually, access nodes consist of multiple linecards and an aggregation card. This structure is shown in Figure 3. While aggregation cards manage high-bandwidth interfaces towards the Broadband Access Servers (BRAS) in the core network, linecards mainly concentrate high numbers of subscribers. Since the paper describes a conceptual framework, the generic term *access node* (AN) is used throughout the paper.

The inherent physical line information, e.g., the port number on the AN, can already be treated as some flavor of LI. Thus, our approach is based on the assumption that LI can be added either by the CPEs (only GPS location information) or by the IPclip mechanism in the ANs (GPS location information *and* access port number *and* access node ID). However, verification and validation of the LI and thereupon taken measures are solely done in the ANs. The reason for doing so is that CPEs are typically not considered as trustworthy network elements by network carriers and service providers. CPEs are usually not within the carriers' management domains. By contrast, ANs are part of the access network and thus within a carrier's management domain. A tuple of information available in ANs is used as precise LI to identify and locate an Internet user:

- the geographic location of the access node
- the access port number the user is connected to
- the access node ID

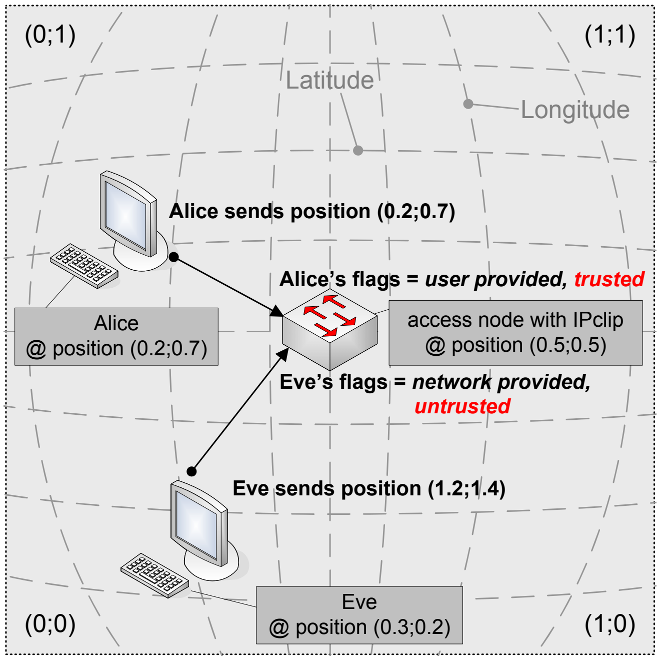
(0;1)
(1;1)
Latitude
Longitude
Alice sends position (0.2;0.7)
Alice's flags = user provided, trusted
Alice
@ position (0.2;0.7)
access node with IPclip
@ position (0.5;0.5)
Eve's flags = network provided, untrusted
Eve sends position (1.2;1.4)
Eve
@ position (0.3;0.2)
(0;0)
(1;0)

Fig. 4. Verification of the location information. The square's edge length (= SCA) and the hosts' positions have been normalized (0 ↔ 1). In a realistic scenario, they are given in GPS coordinates (longitude & latitude) for example.

That is why the IPclip functionality is implemented in the ANs as highlighted in Figure 3.

4.3 Trust Management with IPclip

As already mentioned in the beginning, e-mails shall be fingerprinted and traced by using a *trustworthy* LI. But how can the required level of trustability and credibility be guaranteed? The problem is within the CPEs, which are mostly configured by ordinary Internet users with lack of expertise [38]. Internet users may *unintentionally* misconfigure their CPEs and provide wrong LI. But frauds and scammers may also *intentionally* provide incorrect LI to pretend a false origin of their respective IP traffic. Because of that, CPEs cannot be considered as trustworthy entities. A user provided IPclip option and the LI must thus be verified and checked for plausibility. Additionally, the respective access port number and the access node's ID can be added to the user provided LI.

The IPclip functionality can detect incorrect LI. Therefore, it uses the given fact that only customers can be connected to an AN, which are within a *reasonable geographic distance* relative to that AN. We call this geographic distance the *subscriber catchment area* (SCA) of the respective AN. The SCA is a configurable parameter and defined as the edge length of a square with the AN being located in its center point. To get an idea of the real size of the SCA, the twofold length of the so-called “last mile”, which depends on the deployed transmission technology (DSL, Fiber, Cable), would be a good example. Figure 4 illustrates this setting with normalized coordinates. Two user hosts (Alice and Eve) are located at (0.2;0.7) and (0.3;0.2) respectively. The plausibility of the CPE provided LI is determined by a comparison with the inherent geographic location of the AN (0.5;0.5) with respect to the SCA. If the comparison indicates that IP packets carry incorrect CPE provided LI—maybe due to mobility, misconfiguration, or intentional manipulation—the existing but incorrect LI will be replaced with the inherent LI of the AN. In case that the CPE provides no LI at all, the IPclip mechanism will insert the AN’s inherent geographic location as new IP option with IPclip location information into these IP packets. In any case, IP packets will carry LI about their origin when leaving the IPclip-capable AN—at least with the precision of the AN’s LI and SCA, the port number, and the access node ID or at best with the exact and correct CPE provided LI.

For trust management, special status flags are set during the IPclip verification and validation process. This simple approach differs from typical reputation management systems as reviewed in [39]. These flags give information about the credibility of the LI on the IP level. They are used for control and management but can also be used as triggers for further actions on the application level. Currently, two flags are used, which give conclusions on the origin and on the correctness of the LI. The trust relationship is preserved by these flags at any time since the only places where they are allowed to be assigned are the ANs at the edges of the network carrier’s management domain. And these ANs are under control of the respective carrier or ISP. As a central part of the IPclip system, the naming convention for these flags has been adapted to the commonly used lingo in the area of communication technology. Table 1 briefly summarizes the rough interpretation of the status flags:

`User provided/trusted`: The LI has been provided by the user and has been found correct and plausible during verification (Alice in Figure 4).

`User provided/untrusted`: The LI has been provided by the user. But it did not pass the verification procedures.

`Network provided/untrusted`: The LI has firstly been provided by the user but did not pass verification. Furthermore, the incorrect LI has been overwritten with the AN’s inherent LI (Eve in Figure 4).

`Network provided/trusted`: The IP packets did not carry any user provided LI at all. The LI has been provided by the IPclip mechanism on the AN.

To conclude the overview on the IPclip system, its main tasks are summarized below:

- LI needs to be inserted into *every* IP packet using the IP option format—either by the user/CPE or by IPclip.
- User provided LI must be detected and validated. Existing LI will be overwritten if necessary, e.g., when a user has provided wrong LI.

Table 1. Interpretation of IPclip's Flags

Value	Source / Credibility	Option Description
`00`	`user provided/ untrusted`	A user provided IPclip option did not pass verification.
`01`	`user provided/ trusted`	A user provided IPclip option did pass verification.
`10`	`network provided/ untrusted`	A user provided IPclip option did not pass verification. It is replaced in the AN.
`11`	`network provided/ trusted`	The AN has added a new IPclip option.

- Each AN assigns the status flags inside an IPclip option representing the result of the validation process.
- Moreover, the access port number an the ID of the AN may be added to the IPclip option.
- Optionally, IPclip LI can be removed from IP packets in the downstream data path to not transmit it to the user.

5 IPclip as Anti-Spam Mechanism

Anti-spam mechanisms as revisited in Section 3 try to analyze the content, source & sender information, the reputation, and the behavior of e-mail traffic to derive triggers for e-mail classification and tracing. But all information that is available for analysis is to a very large degree within manipulation-reach of spammers. Hence, typical anti-spam efforts can just react and have to take e-mail traffic as it is. Instead, we try to approach the spam problem from a different perspective. IPclip therefore provides a trustworthy trigger on the IP level, which is *not* within reach of the spammers. As explained in Section 4, the main IPclip system itself is implemented on the ANs and belongs to a network carrier's management domain. The LI and especially the status flags of IPclip are entities, which cannot be manipulated by spammers and can thus be considered as trustable.

The TBW approach and IPclip thereby follow the ideas of [40], i.e. to make ISPs and carriers responsible for the traffic originating in their respective management domains. The primary intention of ISPs should not only be to do *inband traffic control* to protect their own customers because of economic reasons and business models. Instead, ISPs must be encouraged to do *outband traffic control* to increase the overall security level. IPclip is a flavor of outband traffic control.

The questions we try to answer in the following are about how the auxiliary IPclip LI can be utilized in an anti-spam use case, e.g., to trace and/or classify e-mails or even to affect the spammers' elbow-room.

As a starting point, we rely on a prospective setting based on [41] and the following assumptions:

- All ANs within the access network of a self-contained carrier or provider network are equipped with IPclip functionality.
- IPclip handles IP traffic as detailed in Section 4. Thus, when leaving the ANs towards the core network, every IP packet is enriched with an IPclip option with verified LI—either provided by the user's CPE or by the network on the AN. The normal case should be that the CPE provides the LI although this cannot be regarded as mandatory nowadays[3].
- As history has shown, spammers try to disguise their identity and location using different techniques like bot-networks, open relays/proxies, and forged sender information in the e-mail header. By nature, they are not willing to add correct LI to their IP traffic.
- Although differences still exist between the various ISPs, a unified set of "best-practices" in handling e-mails and SMTP must be mandatory. Examples are to enforce SMTP AUTH, to block direct outgoing connections using TCP port 25, and to support sender reputation and e-mail authentication frameworks by default (e.g., SPF, Sender-ID, DKIM). Legal aspects must be regulated on a global scale since e-mailing and the Internet are global as well.

5.1 The Principle Framework

In the anti-spam use case, an important discrepancy needs to be taken into account compared to the general IPclip scenario in Section 4. Originally, IP packets are enriched with LI because IP is a protocol usually providing *direct end-to-end connectivity*. But when doing e-mail via SMTP, IP terminates at each MTA. A new, different IP header is set up when forwarding e-mails to either the next relay, to the receiving mail exchange (MX), or the final recipient. Thus, LI would get lost on the route towards the recipient since *only* the CPEs or the ANs insert IPclip LI at the edges of the network[4]. Hence, LI *must* be preserved—and this must be done on the *first* MTA on the route because the LI is still available on the IP layer. Therefore, we consider the first MTA to save the original LI by copying it as optional, user-defined header fields into the e-mail header. Actually, in doing so, no extra management overhead or reduced performance is expected because an MTA modifies an e-mail's header in any case by inserting its own '`Received: from`' line [4, 41]. Furthermore, the MTA has to gather information from the IP layer, e.g., the source IP address, which is part of this '`Received: from`' line. In Figure 5, we show how these user-defined header fields can look like. The IPclip option as defined in Figure 2 is split up into its individual parts. The responsible MTA additionally specifies its unique identifier. This way, the original IPclip LI, which defines the user host's very first entry point into the carrier network and his geographic location, can be delivered unto the recipient or unto wherever it appears to be useful. As it is a matter of fact

[3] However, suchlike GPS-capable CPE does already exist by now, e.g., the LANCOM 1751 UMTS modem & router with integrated GPS.

[4] Bear in mind that the network edge is the only place where necessary information like a user's access port number is available.

that spammers do forge e-mail headers, they might also forge the IPclip entries. Therefore, a stringent requirement on our framework is that *only* the first MTA on the route is allowed to insert the fields into the e-mail header as Figure 5. How this is assured will be explained in Section 5.2, 5.3 and 5.4.

In the following, we use the term *IP option* when referring to an IPclip option in the IP header. We use the term *SMTP option* when pointing to an IPclip entry in the e-mail header.

```
From - <timestamp>
 X-IPclip-Status: 1100
 X-IPclip-Type: GPS
 X-IPclip-LI: <longitude;latitude>
 X-IPclip-Port: x
 X-IPclip-AN: A
 X-IPclip-MTA: mx.senderhome.net [86.165.10.2]
Return-Path: <sender@senderhome.net>
Received: from ...
```

Fig. 5. Possible structure of an IPclip entry in an e-mail header (= SMTP option) using the format of user-defined fields (`X-...`) as specified in RFC 822.

5.2 A Sender's Options in this Game

E-mail senders are either benign users, direct spammers, or spammer-controlled bots. Spammers show a typical characteristic, which can be exploited using the IPclip mechanisms:

- Spammers lie! Spam pretends to come from everywhere except where it really comes from. Sender information in the header is usually faked and/or the origin of the IP traffic is disguised.
- Good e-mails (ham) do say where they come from. Benign users do not lie.
- In any case, a cautious spammer is not willing to reveal true information about his identity or location.

Relying on these peculiarities of ham and spam, there are four different options in an IPclip-capable network for an e-mail sender in providing LI as IP option. These cases are defined below with respect to IPclip's status flags as summarized in Table 1. Using the status flags as triggers, the *first* MTA might already block incoming e-mails, when they appear untrustworthy. The different possibilities are emblematized in Figure 6. In all four cases, the access node ID (A) and the access port number (x) are added to the LI as well.

1. True location: A user host provides correct LI, which was checked by the AN and validated as `user provided/trusted`. Regarding the IPclip flags,

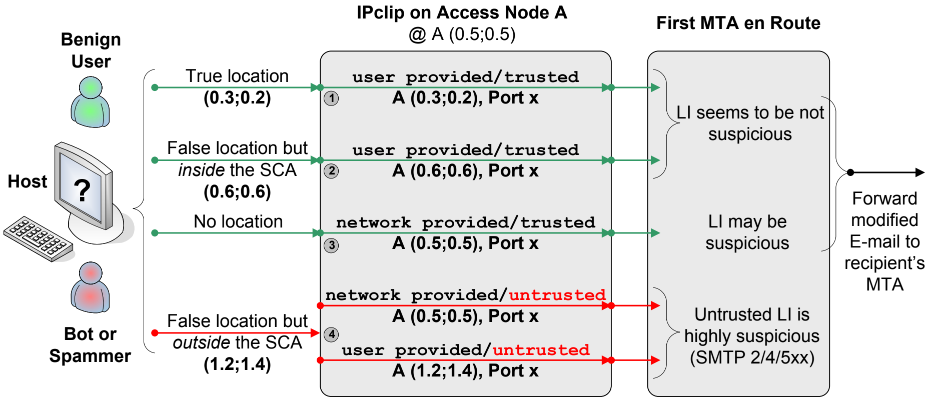

Fig. 6. The different possibilities a host has to insert LI as IP option. A host can either be a benign user, a trojaned bot, or a spammer. The coordinates of the LI have been normalized to 0 ↔ 1 with respect to Figure 4.

the e-mail appears to be trustworthy to the MTA. The IPclip LI is inserted into the e-mail header as proposed in Section 5.1 and the e-mail is forwarded to the next MTA. This way, all e-mails (including spam) would be labeled as `user provided/trusted` but carry inherent location information about the source, which cannot be fiddled any more. Of course, other mechanisms are still required to finally classify an e-mail as spam!

2. False location but inside the SCA: When a user host provides false LI it can nevertheless be labeled as `user provided/trusted` as long as it is *inside* the SCA. Although this appears to be a loophole for spammers, this option only widens the possible geographic area where the sending host is located unto the margins of the AN's respective SCA. Again, the LI is moved into the e-mail header by the MTA before forwarding the modified e-mail to the next hop.

3. No location: When no LI is provided by the CPE, the AN inserts its own LI and sets the flags to `network provided/trusted`. E-mails of this type may be spam since the sending host did not provide any LI at all and should be treated with special care. However, as stated above, we cannot expect every user to have the equipment or the expertise to provide IPclip options. Thus, the MTA modifies the e-mail header accordingly and forwards the e-mail.

4. False location but outside the SCA: In this case, a user provided LI exists but is *outside* the SCA of the AN. It is thus validated as `untrusted`. The IPclip mechanism provides two alternatives: to replace the incorrect LI (`network provided`) or to forward it unchanged (`user provided`). The latter one does not make sense when forwarding the e-mail because the incorrect LI does not contain any usable information. The incorrect LI is thus replaced with the inherent LI of the AN. Now, the first MTA has two choices as well, e.g., depending on local policies. We suggest declining the reception of the e-mail using an SMTP 2/4/5xx return code when it is tagged as `untrusted`. This e-mail is likely to be spam.

Otherwise, when accepting the e-mail, it must be delivered (also from a legal point of view) unless strong evidences justify dropping it. Alternatively, the delivery of suchlike e-mails can be delayed to decrease the e-mail throughput of the untrusted source.

Although spammers do not intend to reveal their true location or identity—neither on SMTP nor on IP level—every IP packet and therewith every e-mail coming from a host is labeled with IPclip LI because an AN *is* required for Internet access. Thereby, the accuracy of the LI is at least the size of the SCA of the respective AN in combination with an access port number.

5.3 Typical E-Mail Flows in Detail

To emblematize our approach, this section describes two typical e-mail routes in detail using Figure 7. The blue path depicts the route of an e-mail from Alice to Bob within the same provider network. The green path highlights the route of an e-mail from Alice to Peter, who is connected to a different provider network. Both providers have a peering/transit agreement to exchange traffic. Alice, Bob, and Peter connect to the Internet via ANs (A, B, C). The provider networks are connected via so-called border gateways (D, E). Basically, these gateways are some special sort of ANs for the respective domain. Both ANs and border gateways are IPclip-capable. Alice's CPE features IPclip as well.

When Alice sends an e-mail to Bob, her IP traffic first passes AN A (1), which checks for existing user provided IP options and verifies them or inserts its own IP option (see Sections 4.3 & 5.2). The e-mail then arrives at Alice's e-mail server MTA1 (2). MTA1 recognizes that it is the first MTA because the IP traffic is enriched with IPclip LI *and* the e-mail does not yet contain an SMTP option[5]. According to Figures 5 and 6, MTA1 inserts the IP option in the e-mail header. If MTA1 is already Bob's MX, the e-mail is directly sent to Bob via AN B (3). Otherwise, MTA1 forwards the e-mail to the respective MX (4), which is MTA2 in this case. MTA2 recognizes that it is not the first MTA because IP options do not exist any more *and* the e-mail does contain the SMTP option. The e-mail finally arrives at Bob (5). By using RF (see Section 4.1) from the IPclip status-field, the SMTP option can either be forwarded to Bob or can be removed by Bob's MX to prevent disclosure of Alice's LI to Bob or due to other on-site or regional policies regarding privacy.

The second mail flow in Figure 7 describes an e-mail's route from Alice to Peter. Steps (1) and (2) are the same as above. Step (6) is similar to (3) and (4) except that the e-mail is not forwarded directly to the recipient but to MTA4—Peter's MX—which belongs to a different provider domain. Border gateway D therefore examines the outbound IP traffic for existing IPclip IP options and inserts its own LI if no option can be detected (status flags set to `network provided/trusted`). This is necessary to remain consistent. Doing so, LI is added to outbound IP traffic again if the original LI has been lost at internal MTAs. If IP options do exists, the traffic is just forwarded because it can

[5] More details about it are given in the summary at the end of this section.

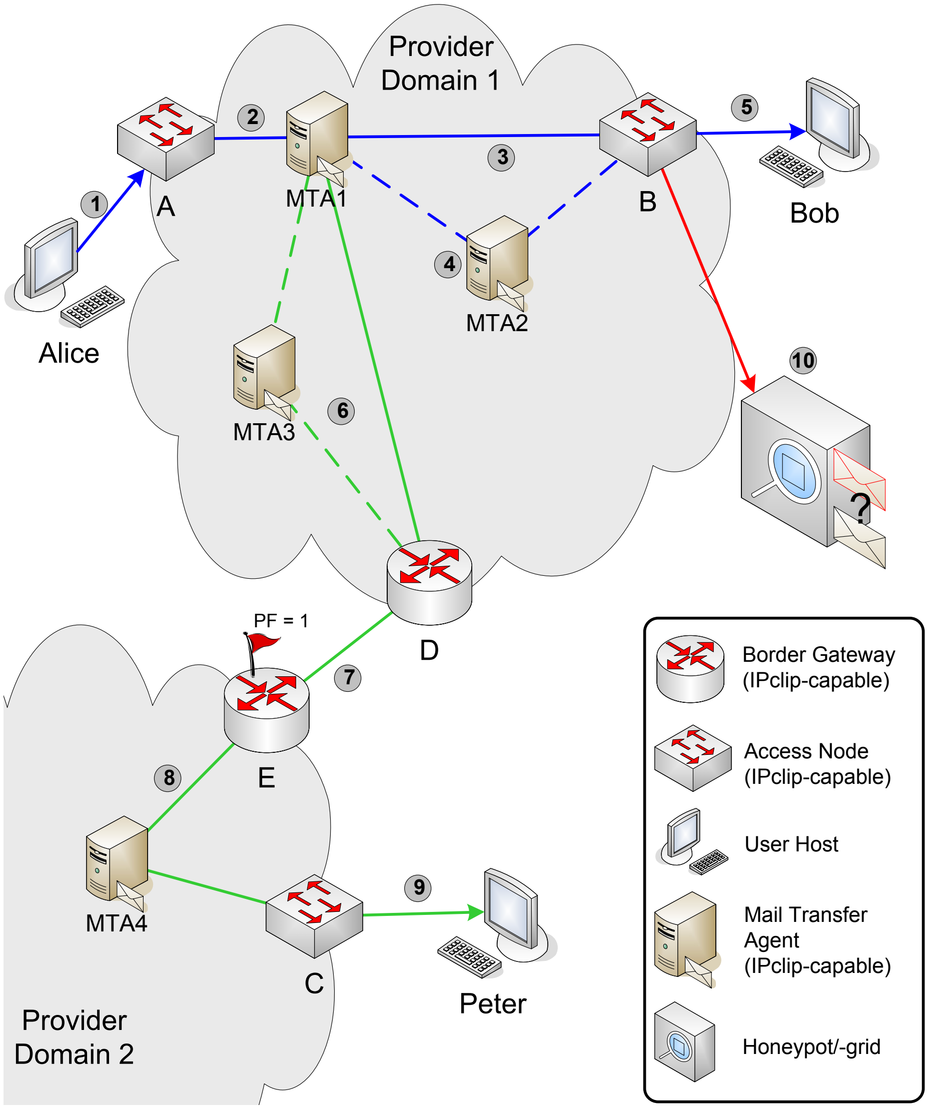

Fig. 7. Two typical mailflows: within the same self-contained network (blue path) and with provider peering (green path). For the latter one, one of the flag inside the IPclip option, the so-called peering flag (PF), is used.

be considered as trustworthy[6]. The border gateways do not examine or touch the e-mail itself. The e-mail arrives at border gateway E of the other provider domain (7). Border gateway E verifies and validates the IP options acting like a normal AN as described in Section 4.3. If the sender's provider did not support IPclip, the gateway's own LI is inserted as IP option. Otherwise, existing IP options are kept and forwarded. In this case, PF (see Section 4.1) within the IPclip status-field is used to indicate that this IP traffic comes from a foreign provider domain. The e-mail is further directed to MTA4 (8). MTA4 recognizes an existing SMTP option and determines an IP option as well. But the operative point is that PF *is set* to `'1'` in the IP option. This way, MTA4 accepts the e-mail without adding or modifying the SMTP option. Without PF set to `'1'`, MTA4 must assume a forged SMTP option. However, the e-mail may lack an SMTP option, e.g., if Alice's provider does not support TBW and IPclip. In that case, MTA4 copies the IP option, which was added by border gateway E, into the header of Alice's e-mail. MTA4 forwards the e-mail to Peter via AN C. Step (9) then equals step (5).

To summarize these two examples, four different cases can be distinguished when an e-mail arrives at an MTA. This number derives from the fact that IPclip LI exists on the IP level and/or on the SMTP level. Thus, there are only four combinations, which represent the base of the proposed framework using the interpretation given below. These cases give clear information about whether this MTA is the first one on the way, an intermediate relay, the recipient's MX, or if an e-mail is not duly transmitted.

Case 1. An IP option but no SMTP option exists: If IPclip LI is present on the IP layer and does not exist in the e-mail header, the traffic can only come from either an AN with PF set to `'0'` or a border gateway with PF set to `'1'`.

If the traffic comes from an AN, a '`Received: from`' line must not yet exist in the e-mail header. The receiving MTA then knows that he is *the first* MTA and inserts the SMTP option in the e-mail header. If a '`Received: from`' line already exists, it must be forged.

Otherwise, if the traffic comes from a border gateway, the MTA is either the first MTA en route (no previous '`Received: from`' line exists) or the first IPclip-capable MTA (at least one '`Received: from`' line exists). An SMTP option is inserted in the e-mail header in either case.

Case 2. Both an IP option and an SMTP option exist: If IPclip LI exists on both layers, PF must be set to `'1'` inside the IP option to indicate that the traffic originated in a different provider domain. Otherwise, something went wrong and the e-mail must be treated with special care.

Case 3. No IP option but an SMTP option exists: If only the SMTP option exists, the MTA is not the first one en route but the receiving MX or and intermediate relay. The e-mail is simply forwarded as usual.

Case 4. Neither an IP option nor an SMTP option does exist: When an e-mail arrives at an MTA lacking IPclip LI on both IP and SMTP layer,

[6] It can only originated from one of the ANs inside the particular sender domain.

something went wrong and this e-mail must be treated with special care. Existing 'Received: from' lines may shed light on the last MTAs that relayed the e-mail though this information should be handled with care since it could be forged as well. Normally, an MTA can figure out the remote station he is speaking SMTP with, e.g, using SPF or (reverse) DNS lookups.

The examples in Figure 7 describe common routes that ham usually takes. But what about the other "good", "ugly", and "bad" mail flows sketched in [13]?

- In a roaming scenario, Alice would connect through an AN and a peering connection to her primary MTA. Mailing lists and forwarding services do not touch the SMTP option, which has been inserted within the primary MTA's domain. These mail flows might have to make use of the peering flag but do not seem to pose a challenge.
- Web-generated e-mails do not have to be spam. Thus, they should always contain an SMTP option with flags set to `user provided/trusted`.
- Outbound connections using TCP port 25 for 'direct-to-MX' delivery or relaying are prohibited by the ISPs by default as assumed for this framework.
- Regarding the problem of open relays and proxies, we still have to rely on up-to-date black lists and the progress made to eliminate these open nodes or to fix misconfigured servers.

5.4 Forging IPclip Fields?

When adding IPclip LI as user-defined fields in e-mail headers, another major question needs to be answered: Is a spammer able to insert forged IPclip fields? In principle, yes! Because everything inside an e-mail header can be forged except the first 'Received: from' line [12]. But falsified 'Received: from' lines and SMTP options can be detected with the proposed framework:

- As we explained, the first MTA knows that it is the first MTA (see Case 1 of the summary above))!
- When arriving at the first MTA, an SMTP option and therewith a forged 'Received: from' line must not exist in an e-mail header.
- The MTA in the first 'Received: from' line must correlate with the MTA named by 'X-IPclip-AN: ...' in the SMTP option (see Figure 6).
- Even if a spammer inserts a forged SMTP option, it has to contain the correct LI of the zombie host he is misusing or at least the LI of the AN this zombie host is connected to including the appropriate access port number. An MTA could compare the potentially forged SMTP option and the original IP option that is validated by the AN. Both need to correlate. Additionally, this would mean to configure individual e-mails for every zombie host resulting in much effort for a spammer.

5.5 Benefits

First and foremost—the proposed framework and IPclip do not recognize spam and do not hinder spam to reach the recipients' inboxes! IPclip is mainly used to recover a useful degree of trustability and unambiguousness (= Trust-by-Wire) since the lack of a trustable, unambiguous reference to the origin of IP traffic is one of the main reasons why spammers can hide as said in the introduction.

1. Tracing Spam: One benefit of the framework is that e-mails can be traced. Actually, tracing and resolving IP addresses is not new, e.g., using WHOIS records [42]. But this has only a coarse grained resolution and is prone to manipulation. Instead, IPclip provides a tuple of information including GPS coordinates, the ID of the AN, and the access port number the sending host is connected to. This appears to be useful since spammers usually hijack hosts that are connected via a flavor of broadband Internet connection. Furthermore, the proposed framework ensures that the LI is trustworthy using various status flags. Thereby, IPclip is independent of dynamic or spoofed IP addresses since the LI refers to the geographic origin and not to a logical address. Classical anti-spam mechanisms have to react on the dynamics and developments of spamming techniques and e-mail forgery [8, 43, 44]. Even legitimate e-mails show different trends regarding the way they are composed and communicated. IPclip is independent of these trends as well.

2. Classifying Spam: A second vantage is that the LI and the status flags can be used as additional triggers in existing anti-spam tools. By allowing only "good" mail flows as described in Section 5.3, some of the loopholes a spammer may exploit are eliminated. A couple of non-ambiguous situations even allow classifying an e-mail as spam. But existing anti-spam tools should remain responsible for the final classification of e-mails. For example, SpamAssassin might add a high penalty score to an e-mail's total scores depending on the status flags of the SMTP option. E-mails, which are labeled as `network provided` and/or `untrusted`, could be delayed for example. Various policies may apply using the IPclip flags as triggers. In the face of high gluts of spam, a *fast* classification of e-mails is still recommended to decide whether the sending host is trustworthy or not or even if it is a benign user or some trojaned bot. Since the IPclip LI is similar to an online equivalent of a caller ID[7], ISPs can more reliably figure out who is sending spam to their customers and more easily block junk e-mail and prosecute spammers. Although the discussed framework envisages the use of additional user-defined header fields, it is compatible to other reputation and sender authentication frameworks. It does neither interfere with SPF nor with DKIM even though DKIM signs selected header fields.

3. Honeynets: A third beneficial aspect is the use of IPclip in combination with honeynets or honeypot MTAs. Honeynets as mentioned in Section 3 are a promising approach to catch and analyze spam e-mails. They can be used to collect information on spam, to extract and analyze IPclip location information, and to trace unsolicited e-mails. As an example, a honeypot host is connected to AN B in Figure 7 analyzing all flavors of e-mails. Thereby, not only e-mails can be traced using the SMTP option but principally the entire IP traffic using the IP option. This way, not only the geographic origin of e-mail traffic can be determined but the origin of remote spamming software as well, which may lead

[7] That is why the name is derived from the ISDN CLIP feature.

to a malware's initial point and initiator. Of course, most traces will not lead to a "big fish" but end at trojaned bots or other loopholes we do not yet know.

5.6 Constraints for the use of IPclip in this Use Case

To guarantee trustworthiness and correct operation of the IPclip mechanism, some requirements and constraints need to be taken into account:

- The existence of an IPclip-capable IP stack is necessary in those network elements, which make use of the IPclip IP option or SMTP option. Other network components do not need to have an IPclip-capable IP stack, since standard-compliant IP options must at least be recognized and skipped using a typical IP-stack. But the complete route through the network must be able to handle IP options.
- A fully IPclip-terminated domain is mandatory. Already a single AN without IPclip functionality discloses a loophole in the network infrastructure similar to an open mail relay or proxy from IPclip's point of view. IP packets with manipulated LI and even fiddled flags can be injected into the network without being validated by a trustworthy IPclip instance. Thus, the presence of IPclip at *all* ANs including border gateways is required. A practicable IPclip domain would be a single self-contained provider or carrier network.
- Legal questions on the availability, the analysis, as well as the storage of sensitive information like the geographic position of Internet users do arise. Actually, the scope of this paper does not primarily cover suchlike privacy issues. Anyway, different possibilities exist to not disclose sensitive LI to the e-mail's recipient.
  1. The IPclip system does not only insert IPclip options and LI into the IP headers in the upstream data path. It is furthermore capable of removing existing IPclip options and LI in the downstream data path if desired.
  2. Using the Removal Flag (RF) in the IPclip status field, the recipient's MX can be configured to delete the LI of the SMTP option in the e-mail header. Thereby, *only* the LI containing the IP packets' geographic origin is removed or replaced with a default value respectively. But the IPclip status field (containing the flags as described in Section 4.1) still remains in the e-mail header to serve as trigger for anti-spam mechanisms at the recipient's place.
  3. Instead of plain text, an encrypted format might be used for the LI inside the IPclip option.

  Moreover, these privacy issues are the same as are already discussed in other areas dealing with the storage of similarly sensitive, private information.

6 Future Work

As the paper described a conceptual framework, future work will cover the following aspects:

- No real life tests have been carried out yet since this would require a fully terminated IPclip domain as mentioned in the previous section. However, a hardware prototype for the IPclip system has been set up using an FPGA development board [45]. Because IPclip is considered to be implemented on access nodes, a hardware solution is necessary to handle IP traffic at line rates of multiple Gbit/s.
- This paper discussed IPclip for IPv4 environments. But IPv6 will be the dominating protocol in the prospective Internet. Thus, IPclip needs to be adapted to IPv6, which provides enhanced mechanisms to convey IP options using extension headers. But as SMTP is above layer 3, these considerations will not touch the proposed anti-spam use case.
- Only e-mail spam has been addressed in this paper. But its pendant in Voice-over-IP (VoIP) calls—so-called *spit* (spam over Internet telephony)—needs to be considered as well. Traditional PSTN services currently migrate into the Internet using different VoIP mechanisms, which tend to supplant or at least measure up with conventional circuit-switched telephony in the future. Telemarketers, prank callers, and other telephone system abusers are likely to target VoIP systems in a similar way as spammers exploit SMTP and e-mail systems. For example, prepared spit hosts are able to call up to thousand VoIP users per minute to play some tape-recorded advertisement. Artless users then call back expensive telephone numbers.

7 Conclusion

The paper reviewed basic aspects and characteristics of spam and the state-of-the-art in anti-spam mechanisms. Afterwards, TBW and IPclip have been revisited in brief. Using this approach, a conceptual anti-spam framework has been proposed. Using IPclip LI as user-defined entries in e-mail headers and guaranteeing the trustability of this LI, e-mails—including spam—can primarily be traced. In some cases, the acceptance can immediately be declined. Furthermore, IPclip provides another trigger for existing anti-spam mechanisms and can be used in honeynets for analysis of IP traffic. Thereby, IPclip approaches the spam problem from a completely different perspective:

- The literal IPclip LI is primarily located on the IP level rather than on SMTP level. It is thus copied to the e-mail header at the first MTA.
- The LI is totally independent from an e-mail's content and SMTP because it refers to the geographic origin of IP traffic. TBW is a flavor of reputation management for the origin of IP traffic. The e-mail header is used as container to convey the LI from the sender to the recipient or to where it might be needed.
- Spammers do not have any influence on the LI and the flags (with respect to the given assumptions and constraints).

Doubtlessly, the discussed framework most likely not considered all possible aspects and loopholes of spammers. But by using GPS as IPclip option type as introduced in Section 4, suchlike LI—either inside the IP header or the e-mail header—points to the origin of IP traffic and e-mails like a finger on a world map. Its existence makes spammers or botnet operators feeling *uncomfortable* because it does not conform to their normal attitude of concealing identity and location.

Acknowledgement

We would like to thank the Broadband Access Division of Nokia Siemens Networks in Greifswald, Germany for their inspiration and continued support in this project. This work is partly granted by Nokia Siemens Networks.

References

[1] L. Kleinrock, “An Internet Vision: The Invisible Global Infrastructure,” *Ad Hoc Networks*, vol. 1, no. 1, pp. 3–11, 2003. 3

[2] ——, “The Internet Rules of Engagement: Then and Now,” *Technology in Society – Technology and Science Entering the 21st Century*, vol. 26, no. 2-3, pp. 193–207, 2004. 3

[3] Jonathan B. Postel, “Simple Mail Transfer Protocol,” RFC 821, August 1982. [Online]. Available: http://tools.ietf.org/html/rfc821 3

[4] J. Klensin, “Simple Mail Transfer Protocol,” RFC 2821, April 2001. [Online]. Available: http://tools.ietf.org/html/rfc2821 3, 13

[5] David H. Crocker, “Standard for the Format of ARPA Internet Text Messages,” RFC 822, August 1982. [Online]. Available: http://tools.ietf.org/html/rfc822 3

[6] D. Streitfeld, “Opening Pandora’s In-Box,” May 2003. [Online]. Available: http://www.latimes.com/technology/la-fi-spam11may11001420,1,5168218,full.story?ctrack=1&cset=true 3

[7] M. Nelson, “Anti-Spam for Businesses and ISPs: Market Size 2003-2008,” Ferris Research, Inc., Tech. Rep., April 2003. [Online]. Available: http://www.ferris.com 3

[8] Symantec Inc., “The State of Spam – A Monthly Report,” Tech. Rep., 2007–2008, appearing monthly. [Online]. Available: http://www.symantec.com/enterprise/security_response/weblog/security_response_blog/spam 3, 4, 20

[9] D. Fallows, “Spam – How It Is Hurting Email and Degrading Life on the Internet,” Ferris Research, Inc., Tech. Rep., October 2003, PEW Internet & American Life Project. [Online]. Available: http://www.pewinternet.org 3

[10] “The Apache SpamAssassin Project.” [Online]. Available: http://spamassassin.apache.org 4, 6

[11] L. H. Gomes, C. Cazita, J. M. Almeida, V. Almeida, and W. M. Jr., “Characterizing a Spam Traffic,” in *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, Taormina, Sicily, Italy, October 2004, pp. 356–369. 4

[12] D. Boneh, “The Difficulties of Tracing Spam Email,” Department of Computer Science, Stanford University, Tech. Rep., September 2004. [Online]. Available: www.ftc.gov/reports/rewardsys/expertrpt_boneh.pdf 5, 19

[13] P. Fältström, “Mail Flows,” Illustration/Artwork, February 2004. [Online]. Available: http://old.openspf.org/mailflows.pdf 5, 19

[14] The Honeynet Project, “Know Your Enemy: Tracking Botnets,” White Paper, March 2005. [Online]. Available: http://www.honeynet.org 5

[15] M. W. Wong, "Sender Authentication – What To Do," White Paper, July 2005, NGS Next Generation Security Software Ltd. [Online]. Available: http://www.openspf.org/blobs/sender-authentication-whitepaper.pdf 5

[16] "Sender Policy Framework Project." [Online]. Available: http://www.openspf.org 5

[17] E. Allman, J. C. M. Delany, M. Libbey, J. Fenton, and M. Thomas, "Domainkeys Identified Mail (DKIM) Signatures," RFC 4871, May 2007. [Online]. Available: http://tools.ietf.org/html/rfc4871 5

[18] M. W. Wong and W. Schlitt, "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1," RFC 4408, April 2006. [Online]. Available: http://tools.ietf.org/html/rfc4408 5

[19] J. Lyon and M. W. Wong, "Sender ID: Authenticating E-Mail," RFC 4406, April 2006. [Online]. Available: http://tools.ietf.org/html/rfc4406 5

[20] "The Spamhaus Project." [Online]. Available: http://www.spamhaus.org 6

[21] J. Carpinter and R. Hunt, "Tightening the Net: A Review of Current and Next Generation Spam Filtering Tools," *Computers & Security*, vol. 25, pp. 566–578, June 2006. 6

[22] "Greylisting.org – A Great Weapon Against Spammers." [Online]. Available: http://www.greylisting.org 6

[23] T. Eggendorfer, "Reducing Spam to 20% of its Original Value with an SMTP Tar Pit Simulator," in *Proceedings of the 2007 MIT Spam Conference on CD-Rom*, Cambridge, MA, USA, March 2007. 6

[24] "419 Eater." [Online]. Available: http://www.419eater.com 6

[25] "Scam o Rama." [Online]. Available: http://www.scamorama.com 6

[26] The Honeynet Project, "Know Your Enemy: Honeynets," White Paper, May 2006. [Online]. Available: http://www.honeynet.org 6

[27] R. Gellens and J. Klensin, "Message Submission," RFC 2476, December 1998. [Online]. Available: http://tools.ietf.org/html/rfc2476 6

[28] ——, "Message Submission for Mail," RFC 4409, April 2006. [Online]. Available: http://tools.ietf.org/html/rfc4409 6

[29] "SpamPlan.com – A simple plan for a spam free life." [Online]. Available: http://www.spamplan.com 6

[30] N. Leavitt, "Vendors Fight Spam's Sudden Rise," *IEEE Computer*, vol. 40, no. 3, pp. 16–19, March 2007. 6

[31] J. Goodman, D. Heckerman, and R. Rounthwaite, "Stopping Spam," *Scientific American*, vol. 292, no. 4, pp. 42–49, April 2005. 6

[32] H. Widiger, S. Kubisch, P. Danielis, J. Schulz, D. Timmermann, D. Duchow, and T. Bahls, "Trust-by-Wire in Packet-switched Networks: Calling Line Identification Presentation for IP," in *1st ITU-T Kaleidoscope Conference – Innovations in NGN*, Geneva, Switzerland, May 2008. 7

[33] S. Kubisch, H. Widiger, P. Danielis, J. Schulz, D. Timmermann, D. Duchow, and T. Bahls, "Countering Phishing Threats with Trust-by-Wire in Packet-switched IP Networks – A Conceptual Framework," in *Proceedings of the 22nd IEEE International Parallel and Distributed Processing Symposium (IPDPS), 4th International Workshop on Security in Systems and Networks (SSN)*, Miami, FL, USA, April 2008. 7

[34] National Marine Electronics Association (NMEA), "NMEA 0183 Standard," January 2002. [Online]. Available: http://www.nmea.de/nmea0183datensaetze.html 7

[35] Information Sciences Institute, University of Southern California, "Internet Protocol Specification," RFC 791, September 1981. [Online]. Available: http://tools.ietf.org/html/rfc791 8

[36] Internet Assigned Numbers Authority, "IP Option Numbers," February 2007. [Online]. Available: http://www.iana.org 8

[37] Agilent Technologies, Inc., "Understanding DSLAM and BRAS Access Devices," White Paper, July 2006. [Online]. Available: http://cp.literature.agilent.com/litweb/pdf/5989-4766EN.pdf 9

[38] S. Gajek, A.-R. Sadeghi, C. Stueble, and M. Winandy, "Compartmented security for browsers – or how to thwart a phisher with trusted computing," in *Proceedings of the 2nd International Conference on Availability, Reliability and Security*, Vienna, Austria, April 2007. 10

[39] S. Ruohomaa, L. Kutvonen, and E. Koutrouli, "Reputation Management Survey," in *Proceedings of the 2nd International Conference on Availability, Reliability and Security*, Vienna, Austria, April 2007. 11

[40] M. Parameswaran, X. Zhao, A. B. Whinston, and F. Fang, "Reengineering the Internet for Better Security," *IEEE Computer*, vol. 40, no. 1, pp. 40–44, January 2007. 12

[41] G. Lindberg, "Anti-Spam Recommendations for SMTP MTAs," RFC 2505, February 1999. [Online]. Available: http://tools.ietf.org/html/rfc2505 12, 13

[42] L. Daigle, "WHOIS Protocol Specification," RFC 3912, September 2004. [Online]. Available: http://tools.ietf.org/html/rfc3912 20

[43] A. Cournane and R. Hunt, "An Analysis of the Tools used for the Generation and Prevention of Spam," *Computers & Security*, vol. 23, no. 2, pp. 154–166, March 2004. 20

[44] J. Graham-Cumming, "The Spammers Compendium," June 2007. [Online]. Available: http://www.jgc.org/tsc.html 20

[45] P. Danielis, S. Kubisch, H. Widiger, J. Schulz, D. Duchow, T. Bahls, D. Timmermann, and C. Lange, "Trust-by-Wire in Packet-switched IP Networks: Calling Line Identification Presentation for IP," in *Design, Automation and Test in Europe Conference and Exhibition (DATE'08), University Booth Hardware Demonstration*, Munich, Germany, March 2008. 22