

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 16, 2007

H. Tschofenig
Nokia Siemens Networks
H. Schulzrinne
Columbia University
D. Wing
J. Rosenberg
Cisco Systems
D. Schwartz
Kayote Networks
June 14, 2007

A Framework for Reducing Spam for Internet Telephony
draft-tschofenig-sipping-framework-sip-reduction-00.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on December 16, 2007.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

Spam, defined as sending unsolicited messages to someone in bulk, might be a problem on SIP open-wide deployed networks. A number of

solutions have been proposed for dealing with Spam for Internet Telephony (SPIT), for example, content filtering, black lists, white lists, consent-based communication, reputation systems, address obfuscation, limited use addresses, turing tests, computational puzzles, payments at risk, circles of trust, and many others. This document describes the big picture that illustrates how the different building blocks fit together and can be deployabled incrementally.

Table of Contents

1.	Introduction	3
2.	Terminology	3
3.	Framework	3
4.	Communication Patterns and User Groups	6
4.1.	Closed Groups	7
4.2.	Semi-Open Groups	7
4.3.	Open Groups	8
4.4.	Summary	8
4.5.	Usability	8
5.	Protocol Interactions	9
5.1.	End Host based Rule Enforcement	9
5.2.	Rule Enforcement via a Trusted Intermediary	9
5.3.	Incremental Deployment	10
6.	Privacy Considerations	12
7.	Example	12
8.	Security Considerations	14
9.	Acknowledgments	15
10.	References	15
10.1.	Normative References	15
10.2.	Informative References	15
Appendix A.	Outside the Scope	17
	Authors' Addresses	17
	Intellectual Property and Copyright Statements	19

[1.](#) Introduction

The problem of Spam for Internet Telephony (SPIT) is an imminent challenge and only the combination of several techniques can provide a framework for dealing with unwanted communication attempts.

[I-D.ietf-sipping-spam] provides four core recommendations that need to be considered for a SPIT solution, namely

- o Strong Identity
- o White Lists
- o Solve the Introduction Problem
- o Don't Wait Until its Too Late

This document illustrates how existing building blocks can be put together to form a solution to deal with SPIT. New building blocks can be added to provide more efficient SPIT handling, since there is no single solution that provides 100 % protection.

The main purpose of this document is largely to define a model of internal device processing, protocol interfaces, and terminology, which define a way in which we can plug-in future protocols. We focus only on the most important solution components and consider many other aspects either outside the scope of this work (see [Appendix A](#)) and postpone them for future work.

[2.](#) Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

[3.](#) Framework

Legend:

oooo: SIP Message Interaction

****: Protocol Interaction for Authorizing the Message Sender

####: Management of Authorization Policies

Figure 1: Overview

Assume that an arbitrary sender transmits a message towards the recipients URI that finally hits the SIP proxy on the recipients side. Information provided within that message are used as input to the rule evaluation. Any part of the message may serve as input to the evaluation process but for practical reasons only a few selected fields do most of the work. In this document, we argue that the authenticated identity of the sender is the most valuable information item. In the future, when standardization progresses then, for example, reputation information obtained from social networks may offer additional input to the authorization process. The protocol

interaction for authorizing the message sender refers to the ability of the recipient or the proxy to interact with the sender to request authorization. The request for authorization is a pull model whereby the proxy or the recipient challenges the sender (e.g., via hash cash [[I-D.jennings-sip-hashcash](#)], or SIP payment [[I-D.jennings-sipping-pay](#)]) for authorization. SIP Identity on the other hand realizes as push model whereby authentication information is pushed towards the recipient.

Figure 2 shows this integration step. The conditions part of the rule offer a mechanisms to incrementally extend the overall framework with new SPIT prevention solution components. Depending on the rule evaluation the message may be rerouted to another entity, such as an answering machine, to the recipient, rejected or other actions are triggered. The latter aspect is particularly interesting since it allows further solution components to be executed. For example, a permission request as part of the consent framework.

SIP message with
authenticated
identity

+-----+

^




Figure 2: Message Filtering and Routing

Note that some traffic analysis and consequently some form of content filtering (e.g., of MESSAGES) can be applied locally within the VoIP provider's domain also under the control of the end user. For example, a user that does not want content filtering to impact message handling would therefore not utilize this functionality.

In a minimalistic SPIT framework only authenticated identities in combination with authorization policies are used. This should serve as a starting point for future work.

Authenticated Identities:

SIP Identity [[RFC4474](#)] is assumed to be used to provide the receiver of a communication attempt with the authenticated identity. SIP Identity is a reasonable simple specification and does not rely on a huge amount of infrastructure support.

Note: SIP Identity is comparable to DKIM

[[I-D.ietf-dkim-overview](#)] used for associating a "responsible" identity with an email message and provides a means of verifying that the association is legitimate.

Authorization Policies:

When the white lists are stored and managed only at the end host then the authorization policies and the protocol to modify the

policies do not need to be standardized since they are purely implementation specific details. Even if the authorization policies are managed centrally or some amount of policy enforcement is done by trusted intermediaries then still there might not be a need to standardize an authorization policy language if the policies can be modified via a webpage. This type of policy handling is done in many cases today already for various applications.

Unfortunately, this approach tends to become cumbersome to manage for end users and therefore it is useful to hide a lot of policy details from the end user itself and to make use of context information, for example, address books and authorization policies available already created for presence based systems.

In some cases the above-described approach is not sufficient whereby it is necessary to define a standardized protocol, for example, if policies are used by different entities, created and modified in an automatic way and when multiple entities manipulate policies (potentially on behalf of the person affected by the policies).

The white list needs to be created somehow and hence there is an introduction problem. [Section 4](#) discusses this aspect in more details.

[4.](#) Communication Patterns and User Groups

When communication takes place then at least three types of groups can be identified.

[4.1.](#) Closed Groups

People in this group communicate only with the peers in their group. They do not appreciate communication attempts from outside. Communication is possible only for people within this list. Here is an example of a closed group: Consider parents that do not want their children from getting contacted by strangers. Hence, they may create a white list containing the identifies of known friends, parents and

other relatives on behalf of their kids.

The usage of authorization policies for usage with Closed Groups is straight forward.

[4.2.](#) Semi-Open Groups

In a semi-open environment all members of the same group are allowed to get in contact with everyone else (e.g., for example in a company environment). For members outside the company the communication patterns depend on the role of the person (e.g., standardization people, sales people, etc.) and on the work style of the person.

For this category we distinguish a number of (non-spam) message sources based on their characteristics:

- o "friends" or "acquaintances", i.e., those we have communicated with before.
- o strangers, divided into 'interesting' and 'uninteresting'. The latter are messages from people that someone does not care to have a conversation with or respond to, at least at that particular moment.

Strangers can be defined by individual names or whole domains. A special class of 'stranger' messages are transaction-related communications, such as Instant Messaging or automated calls from an airline or shipping company.

The usage of authorization policies for usage with Semi-Open Groups can be considered manageable.

In the PSTN a certain amount of protection against unwanted calls is provided due to costs for phone calls. With almost free calls (or instant messages) it might be necessary to abandon the idea of allowing end-to-end real-time message delivery in all cases in order to avoid the alerting the user.

[4.3.](#) Open Groups

People in this type of group are not allowed to limit communication attempts. Help desks, certain people in governmental agencies, banks, insurance companies, etc.

For Open Groups the situation is more complicated. Consider a person working on a customer support helpdesk. Ideally, they would like to receive only calls from friendly customers (although the motivation for calling is most likely a problem) and the topic of the calls only relates to problems they are able to solve. Without listening to the caller they will have a hard time to know whether the call could be classified as SPIT. Many SPIT prevention techniques might not be applicable since blocking callers is likely not possible and applying other techniques, such as turing tests, might not be ideal in an case of Open Groups.

4.4. Summary

Based on the discussions regarding communication patters and groups the following observations can be made:

- o A single person very likely has many roles and they may have an impact on the communication patterns.
For example, consider a person who is working in a company but also want to be available for family members.
- o The context in which a person is may change at any time. For example, a person might be available for family members while at work except during an important meeting where communication attempts may be rejected. Switching a context has an impact for reachability and the means for communicating with a specific recipient, based on enabled rule sets.

From an authorization policy point of view it is important to be able to express a sphere, i.e., the state a user is in and to switch between different spheres easily by thereby switching to a different rule set.

4.5. Usability

An important aspect in the usage of authorization policies is to assist the user when creating policies. Ideally, the policies should be established automatically. Below, there are a couple of examples to illustrate the idea given that these aspects are largely implementation issues:

- o It must be possible for the proxy to automatically add addresses on outbound messages and calls to the rule set. This approach is similar to stateful packet filtering firewalls where outbound packets establish state at the firewall to allow inbound packets to traverse it again.
- o Already available information in the address book can be used for building the policy rules there is quite likely already a relationship available with these persons existent.
- o A large amount of email is non-personal, automated communication, such as newsletters, confirmations and legitimate advertisements. These are often tagged as spam by content filters. This type of correspondence is usually initiated by a transaction over the web, such as a purchase or signing up for a service. [\[I-D.shacham-http-corr-uris\]](#), for example, defines an HTTP header for conveying future correspondence addresses that can be integrated in the rule set.

[5.](#) Protocol Interactions

This section describes the necessary building blocks that are necessary to tie the framework together. We will describe two different environments, namely one where rule enforcement happens at the end host and another one where a trusted network intermediary performs the actions.

[5.1.](#) End Host based Rule Enforcement

- o SIP Identity [\[RFC4474\]](#) is mandatory to implement at the end host and used to determine the authenticated identity of the sending side.
- o Authorization policies are purely implementation specific matter.

Since a purely end host based rule enforcement suffers from a number of drawbacks, rule enforcement by a trusted intermediary is also offered.

[5.2.](#) Rule Enforcement via a Trusted Intermediary

- o SIP Identity [\[RFC4474\]](#) is mandatory to implement at the trusted intermediary (e.g., the immediate VoIP provider) and it determines the authenticated identity of the sending side.
- o Authorization Policies based on the Common Policy framework [\[RFC4745\]](#), as extended in [\[I-D.tschofenig-sipping-spit-policy\]](#) for the purpose of SPIT prevention, are mandatory to implement at the end host side and at the trusted intermediary. The implementation

of the rule evaluation engine might only be necessary on the trusted VoIP proxy. Harmonization with the work done for presence

authorization [[I-D.ietf-simple-presence-rules](#)], which is based on Common Policy [[RFC4745](#)], can be accomplished and is highly desirable.

- o XML Configuration Access Protocol (XCAP) [[I-D.ietf-simple-xcap](#)] is used to create, modify and delete authorization policies and is mandatory to implement at the end host side and at the trusted intermediary.

[5.3.](#) Incremental Deployment

An important property is incremental deployment of additional solution components that can be added and used when they become available. This section aims to illustrate how the extensibility is accomplished, based on an example.

Consider a VoIP provider that provides authorization policies that provide the following functionality equivalent to the Common Policy framework, i.e., identity-based, sphere and validity based conditions initially. For actions only 'redirection' and 'blocking' is provided. In our example we give this basic functionality the AUID 'new-spit-policy-example' with the namespace 'urn:ietf:params:xml:ns:new-spit-policy-example'.

When a client queries the capabilities of a SIP proxy in the VoIP providers network using XCAP the following exchange may take place.

```
GET    /xcap-caps/global/index HTTP/1.1
Host: xcap.example.com
```

Initial XCAP Query for Capabilities

```
HTTP/1.1 200 OK
Etag: "wwhha"
Content-Type: application/xcap-caps+xml
```

```
<?xml version="1.0" encoding="UTF-8"?>
<xcap-caps xmlns="urn:ietf:params:xml:ns:xcap-caps">
```

```

<auids>
  <aid>new-spit-policy-example</aid>
  <aid>xcap-caps</aid>
</auids>
<namespaces>
  <namespace>urn:ietf:params:xml:ns:new-spit-policy-example</namespace>
  <namespace>urn:ietf:params:xml:ns:common-policy</namespace>
</namespaces>
</xcap-caps>

```

Initial XCAP Response with the supported Capabilities

As shown in the example above, Common Policy and the example SPIT extension is implemented and the client can upload rules according to the definition of the rule set functionality.

Later, when the VoIP provider updates the functionality of authorization policies as more sophisticated mechanisms become available and get implemented the functionality of the authorization policy engine is enhanced with, for example, hashcash and the ability to perform statistical analysis of signaling message. The latter functionality comes with the ability to mark messages are Spam and the ability for end users to enable/disable this functionality. We use the AUIDs 'hashcash', 'statistical-analysis' with the namespace 'urn:ietf:params:xml:ns:hashcash' and 'urn:ietf:params:xml:ns:statistical-analysis', respectively.

A end user could now make use of these new functions and a capability query of the SIP proxy would provide the following response.

```

GET    /xcap-caps/global/index HTTP/1.1
Host: xcap.example.com

```

Second XCAP Query for Capabilities

```

HTTP/1.1 200 OK
Etag: "wwhha"
Content-Type: application/xcap-caps+xml

<?xml version="1.0" encoding="UTF-8"?>

```

```

<xcap-caps xmlns="urn:ietf:params:xml:ns:xcap-caps">
  <auids>
    <auid>new-spit-policy-example</auid>
    <auid>xcap-caps</auid>
    <auid>hashcash</auid>
    <auid>statistical-analysis</auid>
  </auids>
  <namespaces>
    <namespace>urn:ietf:params:xml:ns:new-spit-policy-example</namespace>
    <namespace>urn:ietf:params:xml:ns:common-policy</namespace>
    <namespace>urn:ietf:params:xml:ns:hashcash</namespace>
    <namespace>urn:ietf:params:xml:ns:statistical-analysis</namespace>
  </namespaces>
</xcap-caps>

```

Second XCAP Response with the supported Capabilities

New SPIT handling functionality may extend condition, actions and/or transformation elements of a rule.

6. Privacy Considerations

This document does not propose to distribute the user's authorization policies to other VoIP providers nor is the configuration of policies at SIP proxies other than the trusted user's VoIP provider necessary. Furthermore, if blocking or influencing of the message processing is executed by the VoIP provider then they have to be explicitly enabled by the end user. Blocking of messages, even if it is based on "super-clever" machine learning techniques often introduces unpredictability.

Legal norms from fields of law can take regulative effects in the context of SPIT processing, such as constitutional law, data protection law, telecommunication law, teleservices law, criminal law, and possibly administrative law. See, for example, [Law1], [Law2] and [Law3]. For example, it is mandatory to pass full control of SPIT filtering to the end user, as this minimises legal problems.

An overview about regulatory aspects can be found in [Spit-AL].

7. Example

This section shows an example whereby we consider a user Bob@company-example.com that writes (most likely via a nice user interface) the following policies. We use a high-level language to show the main idea of the policies.

RULE 1:

```
IF identity=alice@foo.example.com THEN ACCEPT
IF identity=tony@bar.example.com THEN ACCEPT
```

RULE 2:

```
IF domain=company-example.com THEN ACCEPT
```

RULE 3:

```
IF unauthenticated THEN
    EXECUTE hashcast
    REDIRECT sip:voicebox@company-example.com
```

Example of Bob's Rule Set

At some point in time Bob uploads his policies to the SIP proxy at

Tschofenig, et al.	Expires December 16, 2007	[Page 12]
--------------------	---------------------------	-----------

Internet-Draft	Framework for Reducing Spam	June 2007
----------------	-----------------------------	-----------

his VoIP providers SIP proxy.

```
PUT
/spit-services/users/sip:bob@company-example.com/index/~/spit-services/
service%5b@uri=%22sip:good-friends@example.com%22%5d
HTTP/1.1
Content-Type:application/spit+xml
Host: proxy.home-example.com
```

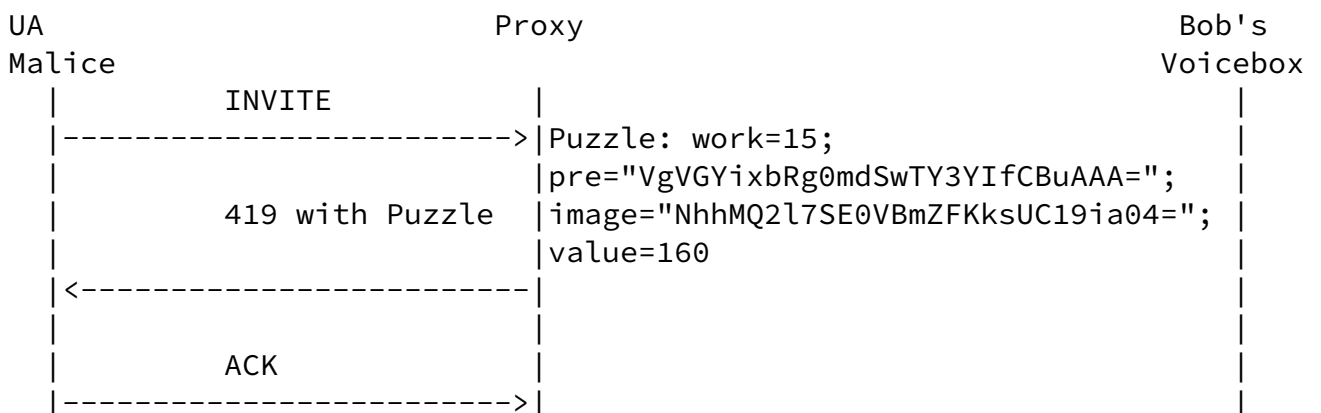
<<<< Policies go in here. >>>>

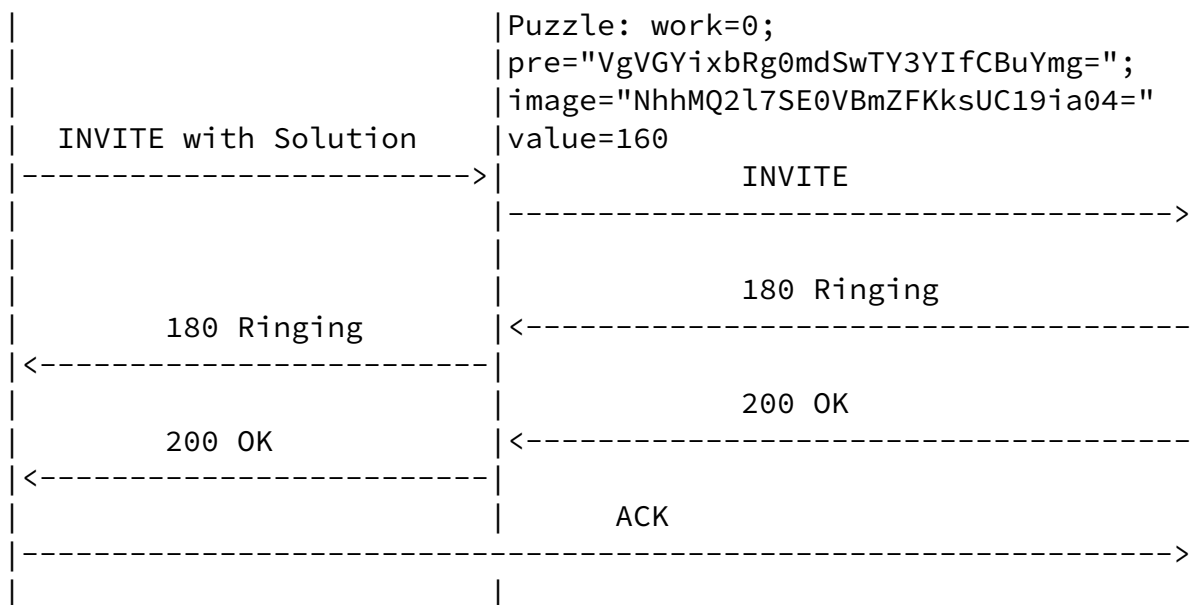
Uploading Policies using XCAP

When BoB receives a call from his friends, alice@foo.example and tony@bar.example.com, then the call is forwarded without any further checks based on Rule 1. The rules assume that the authenticated identity of the caller has been verified.

When Bob receives a call from a co-worker, Charlie@company-example.com, Rule 2 is applied since the domain part in the rule matches the domain part of Charlie's identity.

Now, when Bob receives a contact from an unknown user, called Mallice in this example. Rule 3 indicates that an extended return-routability test using hashcash is used with the call being redirected to Bob's voicebox afterwards. This exchange is shown in the figure below.





Example Exchange: Malice contacts Bob

8. Security Considerations

This document aims to describe a framework for addressing Spam for Internet Telephony (SPIT) in order to make it simple for users to influence the behavior of SIP message routing with an emphasis on SPIT prevention.

The framework relies on three building blocks, namely SIP Identity, authorization policies based on Common Policy and Presence Authorization Policy, and XCAP.

As a high-level overview, the framework allows the user to control end-to-end connectivity at the SIP message routing level whereby the glue that lets all parts fit together is based on authorization policies. Several other solution components can be developed independently and can be plugged into the framework as soon as available.

It must be avoided to introduce Denial of Service attacks against the recipient by misguiding him or her to install authorization policies that allow senders to bypass the policies although that was never intended by the recipient. Additionally, it must not be possible by

extensions to the authorization policy framework to create policies to block legitimate senders or to stall the processing of the authorization policy engine.

9. Acknowledgments

We would like to thank Jeremy Barkan, Dan York, Alexey Melnikov, Thomas Schreck, Eva Leppanen, Cullen Jennings, Marit Hansen, and Markus Hansen for their review comments.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", March 1997.
- [RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", [RFC 4474](#), August 2006.
- [RFC4745] Schulzrinne, H., Tschofenig, H., Morris, J., Cuellar, J., Polk, J., and J. Rosenberg, "Common Policy: A Document Format for Expressing Privacy Preferences", [RFC 4745](#), February 2007.
- [I-D.ietf-simple-xcap] Rosenberg, J., "The Extensible Markup Language (XML) Configuration Access Protocol (XCAP)", [draft-ietf-simple-xcap-12](#) (work in progress), October 2006.

10.2. Informative References

- [I-D.froment-sipping-spit-authz-policies] Froment, T., "Authorization Policies for Preventing SPIT", [draft-froment-sipping-spit-authz-policies-02](#) (work in progress), February 2007.
- [I-D.ietf-sipping-spam] Jennings, C. and J. Rosenberg, "The Session Initiation Protocol (SIP) and Spam", [draft-ietf-sipping-spam-04](#) (work in progress), February 2007.

in progress), February 2007.

[I-D.ietf-simple-presence-rules]

Rosenberg, J., "Presence Authorization Rules",
[draft-ietf-simple-presence-rules-09](#) (work in progress),
March 2007.

[I-D.jennings-sip-hashcash]

Jennings, C., "Computational Puzzles for SPAM Reduction in SIP", [draft-jennings-sip-hashcash-05](#) (work in progress),
June 2007.

[I-D.ietf-sip-consent-framework]

Rosenberg, J., "A Framework for Consent-Based Communications in the Session Initiation Protocol (SIP)",
[draft-ietf-sip-consent-framework-01](#) (work in progress),
November 2006.

[I-D.ietf-dkim-overview]

Hansen, T., Crocker, D., and P. Hallam-Baker, "DomainKeys Identified Mail (DKIM) Message Signing Service Overview",
[draft-ietf-dkim-overview-05](#) (work in progress), June 2007.

[I-D.tschofenig-sipping-spit-policy]

Tschofenig, H., "Anti-SPIT : A Document Format for Expressing Anti-SPIT Authorization Policies",
[draft-tschofenig-sipping-spit-policy-00](#) (work in progress), February 2007.

[I-D.shacham-http-corr-uris]

Shacham, R. and H. Schulzrinne, "HTTP Header for Future Correspondence Addresses", [draft-shacham-http-corr-uris-00](#) (work in progress), May 2007.

[I-D.jennings-sipping-pay]

Jennings, C., "Payment for Services in Session Initiation Protocol (SIP)", [draft-jennings-sipping-pay-05](#) (work in progress), October 2006.

[Spit-AL]

Hansen, M., Hansen, M., Moeller, J., Rohwer, T., Tolkmitt, C., and H. Waack, "Developing a Legally Compliant Reachability Management System as a Countermeasure against SPIT, Third Annual VoIP Security Workshop, Berlin, available at
<https://tepin.aiki.de/blog/uploads/spit-al.pdf>",
June 2006.

Behandlung von Voice over IP (VoIP), available at <http://www.bundesnetzagentur.de/media/archive/3186.pdf>", September 2005.

[Law2] "70. Konferenz der Datenschutzbeauftragten des Bundes und der Laender: Entschliessung Telefonieren mit Internettechnologie (Voice over IP - VoIP), available at http://www.datenschutzzentrum.de/material/themen/press_e/20051028-dsbk-voip.htm", Oktober 2005.

[Law3] "Working Party 29 Opinion 2/2006 on privacy issues related to the provision of email screening services, WP 118, available at http://ec.europa.eu/justice_home/fsj/privacy/docs/wpdocs/2006/wp118_en.pdf", February 2006.

Appendix A. Outside the Scope

We consider the following aspects outside the scope of this document:

- o Mechanisms to publish SPIT causing parties on the Internet, i.e., information about domains that create SPIT.
- o Determining the source of unwanted traffic in real-time.
- o Pushing packet filters and authorization policies towards the SPIT sending domain.

Authors' Addresses

Hannes Tschofenig
Nokia Siemens Networks
Otto-Hahn-Ring 6
Munich, Bavaria 81739
Germany

Email: Hannes.Tschofenig@nsn.com

URI: <http://www.tschofenig.com>

Internet-Draft

Framework for Reducing Spam

June 2007

Henning Schulzrinne
Columbia University
Department of Computer Science
450 Computer Science Building
New York, NY 10027
US

Phone: +1 212 939 7004
Email: hgs+ecrit@cs.columbia.edu
URI: <http://www.cs.columbia.edu>

Dan Wing
Cisco Systems

Phone:
Email: dwing@cisco.com

Jonathan Rosenberg
Cisco Systems
600 Lanidex Plaza
Parsippany, New York 07054
USA

Email: jdrosen@cisco.com
URI: <http://www.jdrosen.net>

David Schwartz
Kayote Networks
Malcha Technology Park

Jerusalem 96951
Israel

Email: david.schwartz@kayote.com

Tschofenig, et al. Expires December 16, 2007 [Page 18]

Internet-Draft Framework for Reducing Spam June 2007

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be

found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).