

SIPPING  
Internet-Draft  
Intended status: Informational  
Expires: January 15, 2009

H. Tschofenig, Ed.  
Nokia Siemens Networks  
G. Dawirs  
University of Namur  
T. Froment  
Alcatel-Lucent  
D. Wing  
Cisco  
H. Schulzrinne  
Columbia University  
July 14, 2008

Requirements for Authorization Policies to tackle Spam and Unwanted  
Communication for Internet Telephony  
draft-froment-sipping-spit-requirements-03

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 January 15, 2009.

Copyright Notice

Copyright (C) The IETF Trust (2008).

## Abstract

Spam over Internet Telephony (SPIT) is one of the foreseen future forms of spamming that SIP open-wide networks may have to handle. SPIT also has more impact on users than email spam since it is more intrusive. Email as a store-and-forward communication mechanism allows for several filtering mechanisms to be applied to the full content before being presented to the user. Session Initiation Protocol (SIP) interaction is, in contrast, real-time communication and therefore does not provide much information prior to the transmission of the content, making it both harder to filter and more annoying to users. The responsibility for filtering, blocking calls, or taking any other preventive action can belong to different elements in the call flow and may depend on various factors. This document discusses the requirements to define authorization policies that should allow end users or other parties to setup anti-SPIT policies for triggering these actions. These policies typically match a particular SIP communication pattern based on a number of attributes. The range of attributes includes information provided, for example, by the SIP protocol itself, by the SIP identity mechanism, by information carried within SAML assertions or by reputation systems of social networks.

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">4</a>
<a href="#">2.</a>	Terminology . . . . .	<a href="#">5</a>
<a href="#">3.</a>	Requirements . . . . .	<a href="#">6</a>
<a href="#">3.1.</a>	Conditions . . . . .	<a href="#">6</a>
<a href="#">3.2.</a>	Actions . . . . .	<a href="#">7</a>
<a href="#">3.3.</a>	Transformations . . . . .	<a href="#">8</a>
<a href="#">3.4.</a>	Generic Requirements . . . . .	<a href="#">8</a>
<a href="#">4.</a>	IANA Considerations . . . . .	<a href="#">9</a>
<a href="#">5.</a>	Security Considerations . . . . .	<a href="#">10</a>
<a href="#">6.</a>	Acknowledgements . . . . .	<a href="#">11</a>
<a href="#">7.</a>	References . . . . .	<a href="#">12</a>
<a href="#">7.1.</a>	Normative References . . . . .	<a href="#">12</a>
<a href="#">7.2.</a>	References . . . . .	<a href="#">12</a>
	Authors' Addresses . . . . .	<a href="#">13</a>
	Intellectual Property and Copyright Statements . . . . .	<a href="#">15</a>

## 1. Introduction

Today, most of the anti-SPAM solutions are coming from email experience, and their applicability to SIP has been discussed in[RFC5039].

As outlined in [[RFC5039](#)], it is likely that many different techniques will need to be combined to deal with SPIT. Users will make different trade-offs when rejecting suspicious calls, for example, trading a lower probability of being interrupted for occasional erroneous call rejection. Also, different types of users, such as businesses and private residences, have different call characteristics. We propose to define a policy language that allows users to easily define their call handling preferences for SPIT. The policy would be executed by trusted SIP proxy or any other SIP element, altering how they handle incoming requests. Policy rules are likely to be define by different actors, including end users themselves, parents on behalf of their children or system administrators. This document enumerates and motivates requirements for such a policy language. Some attributes in an incoming message play a more important role than others. For example, applying authorization policies based on authenticated identity [[RFC4474](#)], is an effective way to make decisions regarding unwanted traffic in some cases.

This document identifies requirements for authorization policies when used to influence message handling for unwanted communication attempts.

## [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](#)], with the important qualification that, unless otherwise stated, these terms apply to the design of the authorization policies, not its implementation or application.

### [3.](#) Requirements

This section lists the requirements categorized according to their applicability for the "conditions", "actions" and "transformation" parts typically found in authorization policies.

#### [3.1.](#) Conditions

The first set of requirements refer to identity related information.

Req-C 1: Policies **MUST** allow conditions to express single authenticated identities.

Req-C 2: Policies **MUST** allow filtering based on the domain part of the identity.

Req-C 3: Policies MUST support the differentiation between authenticated and unauthenticated identities.

Req-C 4: Policies MUST be able to express exceptions within a group of users or a domain.

Req-C 5: Policies SHOULD allow an anonymous identity as a condition.

Message handling may depend on the content of SIP request header fields.

Req-C 6: Policies SHOULD allow conditions to refer to the "destination" (which corresponds to the "Request-URI") and "original-destination" (which corresponds to the "To" header).

Req-C 7: Policies SHOULD allow conditions to refer to the method invoked by the caller (e.g., INVITE, REFER, MESSAGE, SUBSCRIBE).

Motivation: Some SIP methods are more intrusive than others (the default applicative behaviour when SIP MESSAGEs are received is often to pop-up the message on the UAS side), adopting a different filtering policy depending of the method invoked will enhance the user's protection.

Req-C 8: Policies SHOULD allow the entity that writes the rules to take actions on messages that are marked as Spam.

Note that such a condition element should be seen in context of the authenticated domain or, otherwise, of a protected information to avoid security vulnerabilities.

Req-C 9: Policies MAY allow to make decisions based on the current state of the user. E.g., his sphere or other presence information.

Req-C 10: Policies SHOULD support consitions based on the content type and/or offered (or used) media of a message.

Message handling may depend on time of day or the date.

Req-C 11: Policies SHOULD allow conditions that refer to the reception date, time, timezone or period of time of the incoming request.

Message handling might be based on the caller's preferred languages.

Req-C 12: Policies SHOULD allow to make decisions based on the languages in which the originator of the call wishes to communicate.

### [3.2.](#) Actions

Req-A 1: Policies SHOULD allow messages to get "blocked", i.e., to stop forwarding the request and to return an answer with a "403 Forbidden".

Req-A 2: Policies SHOULD allow messages to get "politely blocked", i.e., to stop the request with, for instance, a "486 Busy" response.

Req-A 3: Policies SHOULD allow messages to get "marked", i.e., to forward the request and mark it as "potential Spam" for filtering at the end point or at subsequent entities along the signaling path.

Req-A 4: Policies SHOULD allow messages to be "allowed", i.e., to forward this message.

Req-A 5: Policies MUST allow messages to be "redirected" to, for example, voicemail or to a different device in the possession of the user.

Req-A 6: Policies MUST allow executing other SPIT prevention procedures, such as computational puzzles [[I-D.jennings-sip-hashcash](#)] or the consent framework [[I-D.ietf-sip-consent-framework](#)]. A specification developing a SPIT prevention mechanism should provide information on how they can be incorporated into the authorization policy framework.

Req-A 7: Policies MAY allow an e-mail (or SMS, MMS) or other



notifications to be sent to the user about the actions taken due to a specific call attempt.

R8: Policies MAY allow the usage of one or many feedback mechanisms.

### 3.3. Transformations

Req-T 1: Policies SHOULD allow SIP messages to be marked with a certain SPIT probability in case SPIT detection and policy enforcement is executed on different entities. For example, a network element might run a statistical SPIT detection tool but the authorization policies are executed on a different entity, such as the end host. Note that it needs to be ensured that an adversary is not able to set the SPIT probability values since otherwise the authorization policies that rely on such information are misguided.

### 3.4. Generic Requirements

Req-G 1: It SHOULD be possible to allow a hierarchy of authorization policies to be used.

It is quite likely that a rules from different rule writing entities are provided. For example, in a company environment policies from the system administrator are provided in addition to the end users policies. The former might reflect the overall company policy. The impact for the policy is mainly on the definition of an appropriate conflict resolution mechanism.

Req-G 2: It MUST be possible for a client to learn the supported authorization policy capabilities implemented by the server.

Req-G 3: Policies MUST be extensible and these extensions MUST exist within a different namespace. Furthermore, a published schema and the namespace for elements defined within it MUST NOT be altered by future specifications.

Req-G 4: The policies MUST provide a mandatory-to-implement conflict resolution mechanism.

#### [4.](#) IANA Considerations

This document does not require actions by IANA.

## [5.](#) Security Considerations

This document describes the requirements for elements contained in the authorization policies that allow communication attempts to be treated differently based on the content of the message, time-of-day, context of the user, reputation of the sending party, and many other factors.

The security concerns are related to the ability of certain entities to create, update and delete authorization policies. If an unauthorized entity is allowed to modify policies (and to distribute them to other domains) then a denial of service attack is the consequence with impact for more than a single end point. These security aspects are, however, not the subject of this document.

## 6. Acknowledgements

The content of this document is inspired by the work of CPL [[RFC3880](#)], SIEVE [[RFC5228](#)], Common Policy [[RFC4745](#)] and Presence Authorization Policy [[RFC5025](#)]. We would like to thank the authors of these documents for their work.

Furthermore, we would like to thank Eva Leppanen for the detailed review provided in June 2006.

## [7.](#) References

### [7.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

### [7.2.](#) References

- [I-D.ietf-sip-consent-framework]  
Rosenberg, J., Camarillo, G., and D. Willis, "A Framework for Consent-based Communications in the Session Initiation Protocol (SIP)", [draft-ietf-sip-consent-framework-04](#) (work in progress), January 2008.
- [I-D.jennings-sip-hashcash]  
Jennings, C., "Computational Puzzles for SPAM Reduction in SIP", [draft-jennings-sip-hashcash-06](#) (work in progress), July 2007.
- [RFC3880] Lennox, J., Wu, X., and H. Schulzrinne, "Call Processing Language (CPL): A Language for User Control of Internet Telephony Services", [RFC 3880](#), October 2004.
- [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.
- [RFC5025] Rosenberg, J., "Presence Authorization Rules", [RFC 5025](#), December 2007.
- [RFC5039] Rosenberg, J. and C. Jennings, "The Session Initiation Protocol (SIP) and Spam", [RFC 5039](#), January 2008.
- [RFC5228] Guenther, P. and T. Showalter, "Sieve: An Email Filtering Language", [RFC 5228](#), January 2008.

Tschofenig, et al.

Expires January 15, 2009

[Page 12]

---

Internet-Draft    SPIT Authorization Policy Requirements

July 2008

#### Authors' Addresses

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

Email: [Hannes.Tschofenig@nsn.com](mailto:Hannes.Tschofenig@nsn.com)  
URI: <http://www.tschofenig.com>

Geoffrey Dawirs  
University of Namur  
21, rue Grandgagnage  
Namur B-5000  
Belgique

Email: gdawirs@gdawirs.be

Thomas Froment  
Alcatel-Lucent  
Route de Villejust  
Nozay, Paris 91620  
France

Email: Thomas.Froment@alcatel-lucent.fr

Dan Wing  
Cisco  
170 West Tasman Drive  
San Jose, CA 95134  
USA

Email: dwing@cisco.com

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

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





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](mailto:ietf-ipr@ietf.org).

## Acknowledgment

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