SIPPING
Internet-Draft
Intended status: Informational
Expires: August 30, 2007

G. Dawirs
University of Namur
T. Froment
Alcatel
H. Tschofenig
Siemens

February 26, 2007

Authorization Policies for Preventing SPIT
draft-froment-sipping-spit-authz-policies-02

Status of this Memo

Copyright Notice

Abstract

SPAM, defined as sending unsolicited messages to someone in bulk, might be a problem on SIP open-wide deployed networks. The responsibility for filtering or blocking calls can belong to different elements in the call flow and may depend on various factors. This document discusses mechanisms to establish policies to react on potentially unwanted communication attempts.

These policies match a particular Session Initiation Protocol (SIP) communication pattern based on a number of attributes. The range of attributes includes information provided, for example, by the SIP itself, by the SIP identity mechanism, by information carried within SAML assertions (as introduced with SIP-SAML) and by the SPIT-SAML extensions.

This document raises the question whether it is worth to investigate the aspect of authorization policy usage for SPIT prevention. If so, then the choice of a policy language for describing authorization policies and the details of the authorization policies becomes important. Mechanisms to create, modify and delete authorization policies that are stored in XML documents are already available with XCAP or WEBDAV and they could be reused.

Table of Contents

1. Introduction

The problem of SPAM for VoIP seems to become a very big challenge and only "the combination of several techniques can provide a framework for dealing with spam in SIP" (as stated in [I-D.jennings-sip-hashcash]).

One important building block is to provide a mechanism to instruct some entities in the network to "filter" incoming requests according to user or to network-wide policies. Different entities, such as users or system administrators, might create and modify authorization policies and might even share these policies between domains.

Some attributes in a SIP communication play a more important role than others. For example, applying authorization policies based on the authenticated identity is probably an effective way to accept a communication attempt in order to compat SPIT. The same is true for policies that are applied to deployment friendlier SIP security solutions, such as the SIP identity mechanism [I-D.ietf-sip-identity].

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

The framework of the discussed anti-SPIT authorization policies can be shown as follows:

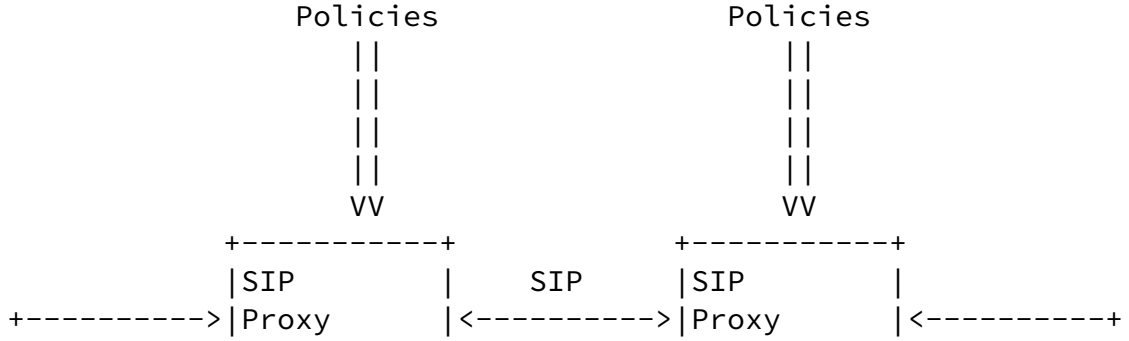

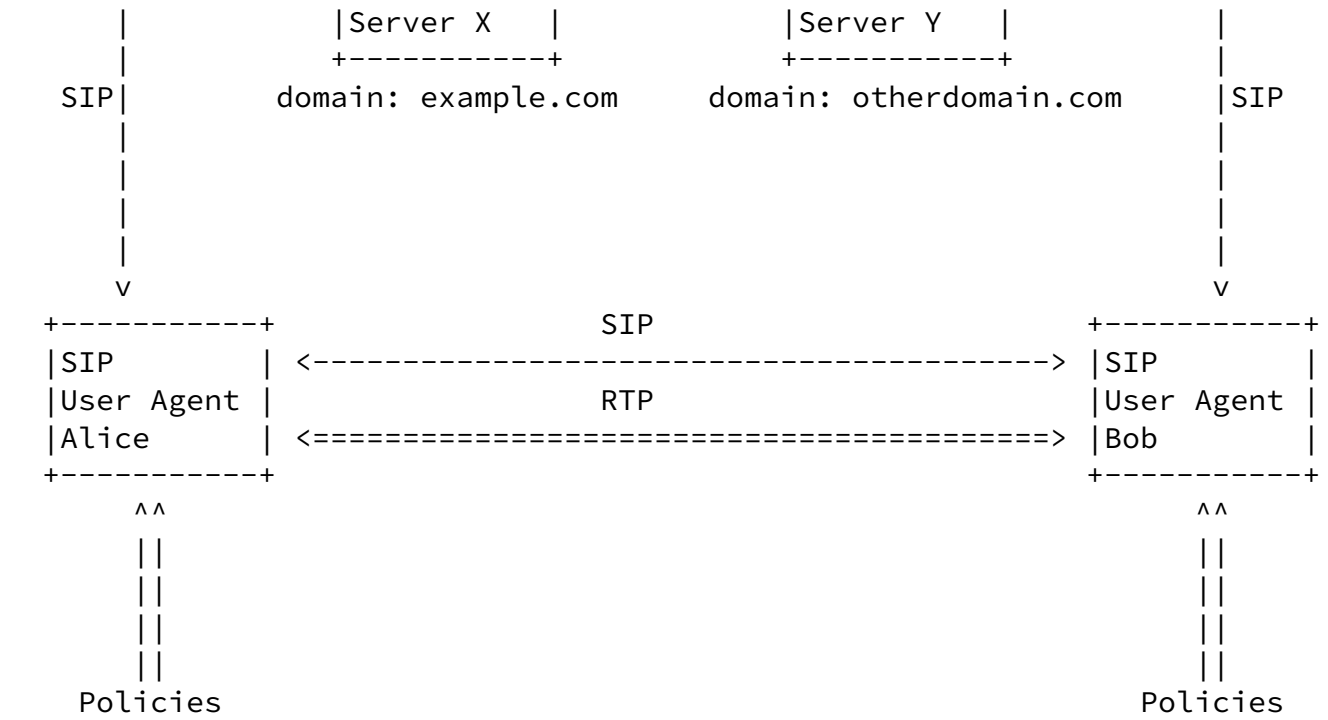

Figure 1: Framework and Scenario

Authorization policies can be applied at the end host and/or by intermediaries. The rule maker might be an end user that owns the device, a VoIP service provider, a person with a relationship to the end user (e.g., the parents of a child using a mobile phone). When the creation, modification and deletion of authorization policies is not only a local matter then a standardized policy language is needed.

The subsequent text lists a few use cases.

The first use case that can be imagined is the case of a user that

asks its outbound proxy to offer protection of requests from a particular SIP UA. He can create an authorization policy rule and upload it to the SIP proxy within its own domain. Requests coming from this SIP URI will then be blocked or treated differently.

Assuming "B@otherdomain.com" has added an authorization rule blocking request coming from domain example.com. Here is a potential message sequence:

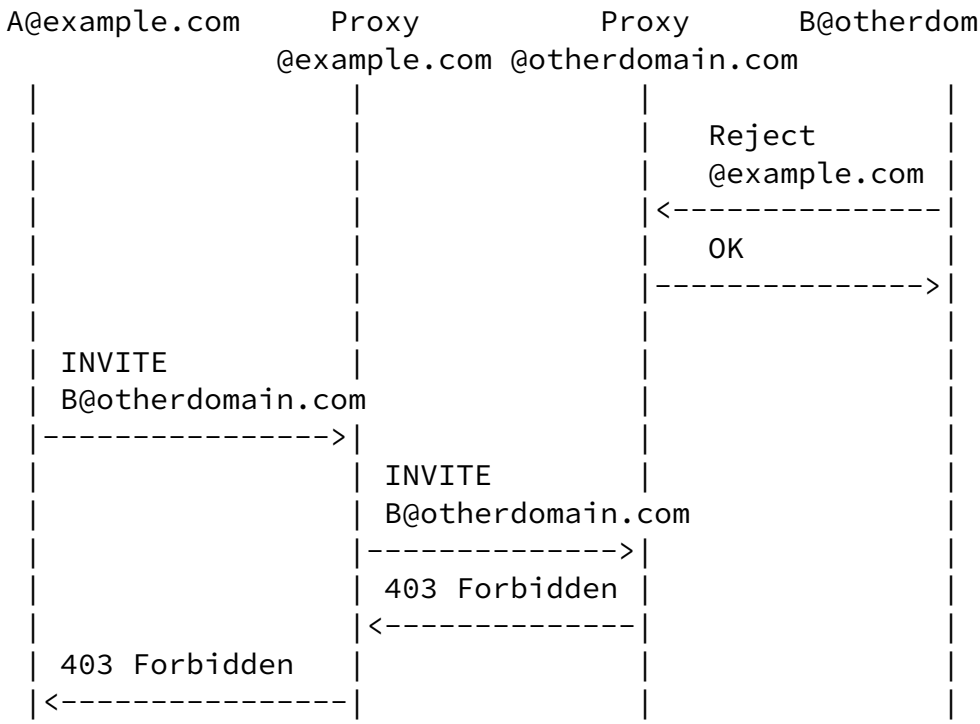

Filtering at the Receiver's Network

If a solution has to be provided to enable SPIT filtering then the following two sub-problems have to be solved:

- An authorization policy language that allows to expression the conditions and actions. An example is [I-D.ietf-simple-presence-rules].

- A mechanism to create, delete, update, retrieve and upload XML based authorization policy rules. XCAP [I-D.ietf-simple-xcap] is a reasonable solution. Another one is WEBDAV [I-D.ietf-webdav-rfc2518bis].

4. Requirements

The design of anti-SPIT authorization policies is guided by the following requirements.

1. The policies SHOULD allow filtering incoming requests depending on several criteria's:

    * Value of any SIP header attribute (e.g., From, To, Contact)

    * Presence (or lack) of any SIP header attribute (e.g., From, To, Contact)

    * Method invoked by the caller (e.g., INVITE, MESSAGE)

    * Value of parameters specified in [I-D.schwartz-sipping-spit-saml]

        + IdentityStrength

        + CostOfCall

        + AuthenticationMethod

        + IdentityAssertion

        + ConnectionSecurity

        + SPITSuspected

        + CallCenter

        + AssertionStrength

    * Request URI of a request

    * Presence of a given expression in the body (subject for further investigation)

2. The policies SHOULD support wildcards (e.g., entire domains)

3. The policies SHOULD support logical operations (and, or, not) between individual elements in conditions

4. The policies SHOULD refer to all authenticated and unauthenticated identities.

5. The policies SHOULD allow a number of actions to be specified, such as:

   * "block": stop forwarding the request and answer with a ``403 Forbidden''

   * "polite-block": drop the request without answering anything

   * "mark": forward the request, putting a flag ``SPAM''

   * "allow": forward this message without conditions (this mechanism is described further)

   * and trigger other mechanism, such as:

     + "puzzle": trigger the "Computational Puzzles" [I-D.jennings-sip-hashcash] mechanism.

     + "consent": trigger the "Consent Framework" [I-D.rosenberg-sipping-consent-framework] mechanism

6. The policies SHOULD allow a default action to be specified.

7. It SHOULD be possible to allow a hierarchy of authorization policies to be used.

5. Discussion and Open Issues

5.1. Extending Geopriv Authorization Policies

The work done in the Geopriv working group on authorization policies seem to be a promising candate for these authorization policies.

To fulfill requirements (1) to (6), it is necessary to decide if [I-D.ietf-geopriv-common-policy] and [I-D.ietf-simple-presence-rules] can be extended.

The following open issues have been identified:

- The authorization policies defined by the Geopriv working group focus on a whitelist approach. This document also raises the question to what extend backlisting capability can be supported or is necessary to support.

- The Geopriv Common Policy mechanism does not allow generic "deny" actions to be defined. This aspect refers to requirements (4) ("all") where (although "all except one" is supported by Common Policy).

- Requirement 2 (wildcards) is provided by Common Policy in a limited fashion by referring to the domain part of an identity. Regular expressions are not supported to keep the policy language simple.

5.2. Hierarchical Authorization Policy Documents

Requirement (7) might require a conflict resolution mechanism to be specified. Geopriv Common Policy currently defines a very simple but powerful conflict resolution mechanism but it is for further investigation whether it is applicable to this problem domain. Other policy languages define a more sophisticated set of conflict resolution mechanisms with preceedence and weights for policies. Although this might be an obviously solution for usage in the context of hierarchical authorization policies it causes problems in other places (such as preserving the order of rules).

6. IANA Considerations

This document does not require actions by IANA.

7. Security Considerations

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.

Furthermore, SPIT prevention techniques often cross the border of what is legally acceptable in certain countries (e.g., filtering SPIT without the consent of the user). Hence, it is extremely important to consider privacy laws in this work.

8. Acknowledgements

Acknowledgements to Yann Lopez for valuable input regarding the usage of Common Policy in the problem domain of SPIT prevention.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", March 1997.

9.2. References

[I-D.ietf-geopriv-common-policy] Schulzrinne, H., "Common Policy: A Document Format for

Expressing Privacy Preferences",
draft-ietf-geopriv-common-policy-11 (work in progress),
August 2006.

[I-D.ietf-simple-presence-rules]
Rosenberg, J., "Presence Authorization Rules",
draft-ietf-simple-presence-rules-08 (work in progress),
October 2006.

[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.

[I-D.ietf-sip-identity]
Peterson, J. and C. Jennings, "Enhancements for
Authenticated Identity Management in the Session
Initiation Protocol (SIP)", draft-ietf-sip-identity-06
(work in progress), October 2005.

[I-D.ietf-webdav-rfc2518bis]
Dusseault, L., "HTTP Extensions for Distributed Authoring
- WebDAV", draft-ietf-webdav-rfc2518bis-18 (work in
progress), February 2007.

[I-D.jennings-sip-hashcash]
Jennings, C., "Computational Puzzles for SPAM Reduction in
SIP", draft-jennings-sip-hashcash-04 (work in progress),
March 2006.

[I-D.rosenberg-sipping-consent-framework]
Rosenberg, J. and J. Rosenberg, "A Framework for Consent-
Based Communications in the Session Initiation Protocol
(SIP)", draft-rosenberg-sipping-consent-framework-00 (work
in progress), July 2004.

[I-D.schwartz-sipping-spit-saml]
Schwartz, D., "SPAM for Internet Telephony (SPIT)
Prevention using the Security Assertion Markup Language
(SAML)", draft-schwartz-sipping-spit-saml-01 (work in

progress), June 2006.

Appendix A. Sophisticated SPIT Filtering Scenario

In a more sophisticated scenario one might even consider the idea of stopping SPIT as early as possible. Domains may agree to exchange authorization policies in order to stop SPIT earlier (i.e., closer to the source of the problem). The subsequent text describes this scenario.

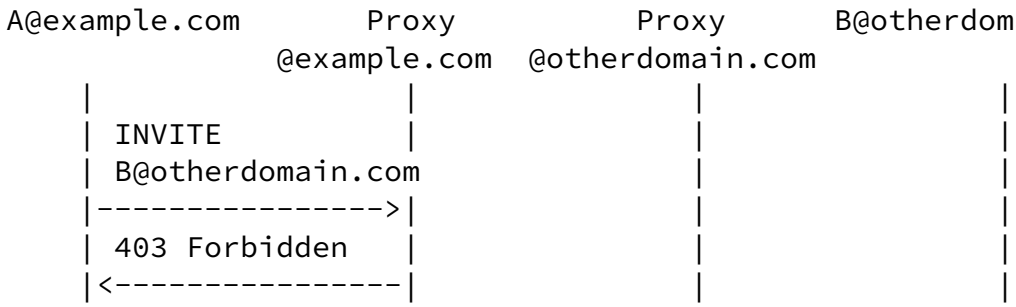

Filtering at the Sender's Network

This call flow illustrates the bandwidth-saving interest of this use case.

Though, two scenarios could happen:

- In the good case, the sender's domain is honest and exchanges authorization policies in order to apply rules that avoids forwarding unsollicited requests.

- In the worst case, the sender's domain is not cooperative. It will refuse to upload such documents. In this case, the presence of rules in the recipient's domain will suffice to keep the recipient "SPAM free", even if more traffic has been consumed (since the request has been relayed at least until the first proxy of the recipient's domain, exactly like in the first use case).

It might be desirable to use a hierarchy of authorization policy documents that need to be combined when applying them to the SIP signaling traffic. This raises the question of a merging algorithm, particularly when authorization policy rules are conflicting or contain blacklists.

This scenario is subject for further discussion since it might raise privacy concerns when privacy policies are shared and potentially applied to other communication partners traffic.

Authors' Addresses

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

Email: gdawirs@gdawirs.be

Thomas Froment
Alcatel
1, rue Ampere - BP 80056
Massy, Paris 91302
France

Email: Thomas.Froment@alcatel-lucent.fr

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

Email: Hannes.Tschofenig@siemens.com
URI: http://www.tschofenig.com

Full Copyright Statement

Intellectual Property

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).