Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 30, 2007

H. Tschofenig
Siemens Networks GmbH & Co KG
D. Wing
Cisco
H. Schulzrinne
Columbia U.
T. Froment
Alcatel-Lucent
G. Dawirs
University of Namur
February 26, 2007

Anti-SPIT : A Document Format for Expressing Anti-SPIT Authorization Policies

draft-tschofenig-sipping-spit-policy-00.txt

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 defines an authorization based policy language that allows end users to upload anti-SPIT policies to intermediaries, such as SIP proxies. These policies mitigate unwanted SIP communications. It extends the Common Policy authorization framework with additional conditions and actions. The new conditions 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 SIP itself, by the SIP identity mechanism, by information carried within SAML assertions.

Table of Contents

1. Introduction 4
2. Terminology 4
3. Generic Processing 5
3.1. Structure of SPIT Authorization Documents 5
3.2. Rule Transport 5
4. Condition Elements 5
4.1. MessagePattern Element 6
4.2. MethodUsed Element 6
4.3. Assertions-Specific Parameters 6
5. Actions 7
5.1. Handling Action 7
5.2. Redirect Action 8
6. Examples 8
7. XML Schema 11
8. XCAP USAGE 12
8.1. Application Unique ID 12
8.2. XML Schema 12
8.3. Default Namespace 12
8.4. MIME Type 13
8.5. Validation Constraints 13
8.6. Data Semantics 13
8.7. Naming Conventions 13
8.8. Resource Interdependencies 13
8.9. Authorization Policies 13
9. IANA Considerations 13
9.1. Anti-SPIT Policy XML Schema Registration 13
9.2. Anti-SPIT Policy Namespace Registration 14
9.3. XCAP Application Usage ID 14
10. Security Considerations 14
11. Contributors 15
12. Acknowledgments 15
13. References 15
13.1. Normative References 15
13.2. Informative References 15

Authors' Addresses 16
Intellectual Property and Copyright Statements 18

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, as stated in [11].

One important building block is to have a mechanism that can instruct SIP intermediaries to react differently on incoming requests based on policies. Different entities, such as end users, parents on behalf of their children, system administrators in enterprise networks, etc., might create and modify authorization policies. The conditions in these policies can be created from many sources but some information elements are more important than others. For example, there is reason to believe that applying authorization policies based on the authenticated identity is an effective way to accept a communication attempt to deal with unsolicited communication. Authentication based on the SIP identity mechanism, see [2], is one important concept.

There is also related work in this context that needs to be highlighted. Requirements for the authorization policies described in this document are outlined in [7]. Selected parts of the work done with Sieve [13], a mail filtering language, may be reused by this document. Furthermore, the Call Processing Language (CPL) [14] is similar to the approach described in this document. The difference mainly is that CPL has a more procedural approach, while this proposal is matching-based.

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 [1].

This document reuses the terminology from RFC 4745 [3]:

Rule maker:

> The RM is an entity that creates the authorization policies that react to unwanted connection attempts. 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). A standardized policy language is needed when the creation, modification and deletion of authorization policies are not only a local matter.

Authorization policy:

> An authorization policy is given by a rule set. A rule set contains an unordered list of rules. Each rule has a condition, an action and a transformation component. The terms 'authorization policy', 'policy', 'rule set', 'authorization policy rule', 'policy rule' and 'rule' are used interchangeably. Authorization policies can be applied at the end host and/or by intermediaries.

Permission:

> The term permission refers to the action and transformation components of a rule.

3. Generic Processing

3.1. Structure of SPIT Authorization Documents

A SPIT authorization document is an XML document, formatted according

to the schema defined in RFC 4745 [3]. SPIT authorization documents inherit the MIME type of common policy documents, application/ auth-policy+xml. As described in [3], this document is composed of rules which contain three parts - conditions, actions, and transformations. Each action or transformation, which is also called a permission, has the property of being a positive grant to the authorization server to perform the resulting actions, be it allow, block etc . As a result, there is a well-defined mechanism for combining actions and transformations obtained from several sources. This mechanism therefore can be used to filter connection attempts thus leading to effective SPIT prevention.

3.2. Rule Transport

Policies are XML documents that are stored at a Proxy Server or a dedicated device. The Rule Maker therefore needs to use a protocol to create, modify and delete the authorization policies defined in this document. Such a protocol is available with the Extensible Markup Language (XML) Configuration Access Protocol (XCAP) [4].

4. Condition Elements

This section describes the additional enhancements of the conditions-part of the rule. This document inherits the Common Policy functionality, including identity, validity, and sphere conditions.

The identity condition restricts matching of a rule either to a single entity or a group of entities. Authenticated and non-authenticated entities can be matched; acceptable means of authentication are specified in Section 3.1 of [10] and can be reused in this document. An important component of the overall solution are authenticated identities, such as provided via SIP Identity [2]. If the <identity> element is absent, identities are not considered, and thus, other conditions in the rule apply to any user, authenticated or not.

The <identity> condition is considered TRUE if any of its child elements (e.g., the <one> and the <many> elements defined in this document) evaluate to TRUE, i.e., the results of the individual child element are combined using a logical OR.

4.1. MessagePattern Element

Any attribute of the SIP header, such as the From, To, Contact etc., can be used to perform actions on incoming messages.

4.2. MethodUsed Element

Any SIP Method invoked by the user can be used to filter incoming messages.

4.3. Assertions-Specific Parameters

This parameter list set refers to information that can be made available by, for example, using SAML assertions, as defined in [5]. As an example, the following attribute is reused in this document:

AuthenticationOfAccountOpening:

(a) No validation of new account (could be machine opened)
(b) Turing Test (human needed to open new account)
(c) Credit card or other form of verifiable identification
(d) Passport was presented for verification

The values put in the element are defined as follows:

- Corresponds to value (a)- (d)
- Corresponds to value (b) - (d)
- Corresponds to value (c) - (d)

- Corresponds to value (d)

Other attributes, such as IdentityStrength, CostOfCall, IdentityAssertion, ConnectionSecurity, SPITSuspected, CallCenter, or AssertionStrength from [5] might allow meaningful decisions to be performed.

Further parameters carried in a SAML assertion are defined in [6] and can also be used for the decision making process. Possible parameters for Originating Line Indication (OLI) and for Calling Party Category (CPC) are described in Section 7 and 8 of [6]. The CPC parameters may also be encoded in a different form, as shown in [8], and usable by this document.

5. Actions

As stated in [2], conditions are the 'if'-part of rules, whereas actions and transformations form their 'then'-part. The actions and transformations parts of a rule determine which operations the proxy server MUST execute on receiving a connection request attempt that matches all conditions of this rule. Actions and transformations permit certain operations such as block, polite-block, mark, allow, puzzle and consent.

5.1. Handling Action

The <handling> element allows a couple of actions to be defined.

Block Action:

The block action states that this specific connection request MUST NOT be forwarded and a "403" forbidden message MUST be sent to the sender of the message.

Polite-block Action

The Polite-block action states that this specific connection request MUST NOT be forwarded and no message be sent back to the sender of the message.

Mark Action:

The Mark action states that this specific connection request MUST be forwarded after marking it as a "SPAM". Details for the message marking are for further study.

Allow Action:

The Allow action states that this specific connection request MUST be forwarded.

Puzzle Action:

The Puzzle action states that the "Computational Puzzles" mechanism, described in [11], MUST be triggered.

Consent Action:

The Consent action states that "Consent Framework" [12] mechanism MUST be triggered.

Default Action:

One of the action can be stated as a default action.

5.2. Redirect Action

This document defines the <redirect> action that contains a URI where an incoming message is forwarded to.

6. Examples

This section provides a few examples for policy rules defined in this document. The example policy shows three rules with the rule id 1, 2 and 3. The rule with the id=1 matches for authenticated identities from the domain "example.com", "example.org" and the single identity "sip:bob@good.example.net". For these conditions SIP messages are forwarded to the SIP UA as indicated with the <handling> element.

Rule 2 indicates that for SIP messages where the identity cannot be matched against a white list and for those where the identity was obtained by having the user to present a passport, credit card or other form of verifiable identification when opening the account (as indicated in the <AuthenticationOfAccountOpening> by setting the token 'VERIFYABLE') the consent framework is applied (see 'consent' token in the <handling> element).

Rule 1 and 2 are valid only from 2007-1-24T17:00:00+01:00 to 2007-3-24T19:00:00+01:00.

Rule 3 does not contain any condition. All requests that fall into this category are redirected to an answering machine (namely sip:answering-machine@home.foo-bar.com). Rule 3 is not restricted to

a specific time period.

```
<?xml version="1.0" encoding="UTF-8"?>
<ruleset xmlns="urn:ietf:params:xml:ns:common-policy"
  xmlns:spit="urn:ietf:params:xml:ns:spit-policy"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">

  <rule id="1">
    <conditions>
      <identity>
        <one id="sip:bob@good.example.net"/>
        <many domain="example.com"/>
        <many domain="example.org"/>
      </identity>
      <validity>
        <from>2007-1-24T17:00:00+01:00</from>
        <until>2007-3-24T19:00:00+01:00</until>
      </validity>
    </conditions>
    <actions>
      <spit:handling>allow</spit:handling>
    </actions>
    <transformations/>
  </rule>

  <rule id="2">
    <conditions>
      <validity>
        <from>2007-1-24T17:00:00+01:00</from>
        <until>2007-3-24T19:00:00+01:00</until>
      </validity>
      <spit:AuthenticationOfAccountOpening>VERIFYABLE
      </spit:AuthenticationOfAccountOpening>
    </conditions>
    <actions>
      <spit:handling>consent</spit:handling>
    </actions>
    <transformations/>
  </rule>

  <rule id="3">
```

```
      <conditions/>
      <actions>
        <spit:redirect>sip:answering-machine@home.foo-bar.com
        </spit:redirect>
      </actions>
      <transformations/>
    </rule>

  </ruleset>
```


7. XML Schema

This section contains the XML schema that defines the policies schema described in this document. This schema extends the Common Policy schema (see [2]) by introducing new members of the <condition> and <action> elements.

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema
  targetNamespace="urn:ietf:params:xml:ns:spit-policy"
  xmlns:spit="urn:ietf:params:xml:ns:spit-policy"
  xmlns:cp="urn:ietf:params:xml:ns:common-policy"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  elementFormDefault="qualified"
  attributeFormDefault="unqualified">

  <!-- This import brings in the XML language attribute xml:lang-->
  <xs:import namespace="http://www.w3.org/XML/1998/namespace"
    schemaLocation="http://www.w3.org/2001/xml.xsd"/>

  <!-- Conditions -->

  <xs:element name="MethodUsed"
    type="xs:string"/>

  <xs:element name="CostOfCall"
    type="xs:integer" default="0"/>

  <xs:element name="IdentityStrength"
    type="xs:integer" default="0"/>
```

```
    <xs:element name="AuthenticationOfAccountOpening">

      <xs:simpleType>
        <xs:restriction base="xs:token">
          <xs:enumeration value="NO_EFFORT"/>
          <xs:enumeration value="HUMAN"/>
          <xs:enumeration value="VERIFYABLE"/>
          <xs:enumeration value="PASSPORT"/>
        </xs:restriction>
      </xs:simpleType>
    </xs:element>

    <xs:element name="MessagePattern">
      <xs:complexType>
          <xs:attribute name="context" />
      </xs:complexType>
```



```
    </xs:element>

    <!-- Action -->

    <xs:element name="handling">
      <xs:simpleType>
        <xs:restriction base="xs:token">
          <xs:enumeration value="block"/>
          <xs:enumeration value="mark"/>
          <xs:enumeration value="polite-block"/>
          <xs:enumeration value="allow"/>
          <xs:enumeration value="puzzle"/>
          <xs:enumeration value="consent"/>
        </xs:restriction>
      </xs:simpleType>
    </xs:element>

    <xs:element name="redirect" type="xs:string" />

  </xs:schema>
```

8. XCAP USAGE

The following section defines the details necessary for clients to

manipulate SPIT authorization documents from a server using XCAP.

8.1. Application Unique ID

XCAP requires application usages to define a unique application usage ID (AUID) in either the IETF tree or a vendor tree. This specification defines the "Spit-policy" AUID within the IETF tree, via the IANA registration in Section 9.

8.2. XML Schema

XCAP requires application usages to define a schema for their documents. The schema for Anti-SPIT authorization documents is described in Section 7.

8.3. Default Namespace

XCAP requires application usages to define the default namespace for their documents. The default namespace is urn:ietf:params:xml:ns:spit-policy.

8.4. MIME Type

XCAP requires application usages to defined the MIME type for documents they carry. Anti-SPIT privacy authorization documents inherit the MIME type of Common Policy documents, application/ auth-policy+xml.

8.5. Validation Constraints

This specification does not define additional constraints.

8.6. Data Semantics

This document discusses the semantics of Anti-SPIT authorization.

8.7. Naming Conventions

When a SIP Proxy receives a SIP message to route it towards to a

specific user foo, it will look for all documents within http://[xcaproot]/spit-policy/users/foo, and use all documents found beneath that point to guide authorization policy.

8.8. Resource Interdependencies

This application usage does not define additional resource interdependencies.

8.9. Authorization Policies

This application usage does not modify the default XCAP authorization policy, which is that only a user can read, write or modify his/her own documents. A server can allow privileged users to modify documents that they do not own, but the establishment and indication of such policies is outside the scope of this document.

9. IANA Considerations

There are several IANA considerations associated with this specification.

9.1. Anti-SPIT Policy XML Schema Registration

URI: urn:ietf:params:xml:schema:spit-policy

Registrant Contact: Hannes Tschofenig (hannes.tschofenig@siemens.com).

XML: The XML schema to be registered is contained in Section 7. Its first line is

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

and its last line is

```
</xs:schema>
```

9.2. Anti-SPIT Policy Namespace Registration

URI: urn:ietf:params:xml:ns:spit-policy
Registrant Contact: Hannes Tschofenig
(hannes.tschofenig@siemens.com).
XML:

9.3. XCAP Application Usage ID

This section registers an XCAP Application Usage ID (AUID) according to the IANA procedures defined in .

Name of the AUID: spit-policy

Description: Anti-SPIT privacy rules are documents that describe the Authorization policies that trigger reaction to unwanted connection attempts.

10. Security Considerations

This document aims to make it simple for users to influence the behavior of SIP message routing with an emphasis on SPIT prevention. This document proposes a strawman proposal for conditions and actions that might be useful when it comes to allowing a UA to tell its proxies which messages it wants to receive and what tasks it wants those proxies to perform before sending a SIP request to the UA.

A couple of requirements are described in [7] and a general discussion about the available solution mechanisms is available with [9]. This document offers the ability to glue the different solution pieces together.

Since this document uses the Common Policy framework it also inherits its capabilities, including the combining permission algorithm that is applied when multiple rules fire. Unauthorized access to the user's Anti-SPIT rules must be prevented to avoid the introduction of

security vulnerabilities.

11. Contributors

We would like to thank Mayutan Arumaithurai (mayutan.arumaithurai@gmail.com) for his work on this document.

12. Acknowledgments

We would like to thank David Schwartz for his work on the "SAML SPIT" draft. We would like to thank Miguel Garcia and Remi Denis-Courmont for their review comments.

Finally, we would like to thank Jonathan Rosenberg, David Schwartz and Dan York for sharing their thoughts with us.

13. References

13.1. Normative References

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

[2] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, August 2006.

[3] 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.

[4] Rosenberg, J., "The Extensible Markup Language (XML) Configuration Access Protocol (XCAP)", draft-ietf-simple-xcap-12 (work in progress), October 2006.

13.2. Informative References

[5] 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.

[6] Schubert, S., "Conveying CPC using the SAML", draft-schubert-sipping-saml-cpc-02 (work in progress), July 2006.

[7] Froment, T., "Authorization Policies for Preventing SPIT", draft-froment-sipping-spit-authz-policies-01 (work in progress), June 2006.

[8] Mahy, R., "The Calling Party's Category tel URI Parameter", draft-mahy-iptel-cpc-05 (work in progress), October 2006.

[9] Jennings, C. and J. Rosenberg, "The Session Initiation Protocol (SIP) and Spam", draft-ietf-sipping-spam-03 (work in progress), October 2006.

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

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

[12] 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.

[13] Showalter, T., "Sieve: A Mail Filtering Language", RFC 3028, January 2001.

[14] Lennox, J., Wu, X., and H. Schulzrinne, "Call Processing Language (CPL): A Language for User Control of Internet Telephony Services", RFC 3880, October 2004.

Authors' Addresses

Hannes Tschofenig
Siemens Networks GmbH & Co KG
Otto-Hahn-Ring 6
Munich, Bavaria 81739
Germany

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

Dan Wing
Cisco

Phone:
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+ecrit@cs.columbia.edu
URI: http://www.cs.columbia.edu

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

Email: Thomas.Froment@alcatel-lucent.fr

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

Email: gdawirs@gdawirs.be

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).