SIP
Internet-Draft
Intended status: Informational
Expires: May 3, 2009

M. Munakata
S. Schubert
T. Ohba
NTT

October 30, 2008

UA-Driven Privacy Mechanism for SIP
draft-ietf-sip-ua-privacy-03

Status of this Memo

Abstract

This document defines a best current practice for a user agent to generate an anonymous SIP message by utilizing mechanisms such as GRUU (Globally Routable User Agent URIs) and TURN (Traversal Using Relays around NAT) without the need for a privacy service defined in RFC 3323.

Table of Contents

1. Introduction 3
2. Terminology 3
3. Concept of Privacy 3
4. Treatment of Privacy-Sensitive Information 4
4.1. Obtaining a Functional Anonymous URI Using the GRUU Mechanism 4
4.2. Obtaining a Functional Anonymous IP Address Using the TURN Mechanism 5
5. User Agent Behavior 6
5.1. Critical Privacy-Sensitive Information 6
5.1.1. Contact Header Field 6
5.1.2. From Header Field 7
5.1.3. Via Header Field 8
5.1.4. IP Addresses in SDP 8
5.2. Non-Critical Privacy-Sensitive Information 8
5.2.1. Host Names in Other SIP Header Fields 8
5.2.2. Optional SIP Header Fields 8
6. Security Considerations 9
7. IANA Considerations 9
8. References 9
8.1. Normative References 9
8.2. Informative References 10
Authors' Addresses 10
Intellectual Property and Copyright Statements 12

1. Introduction

[RFC3323] defines a privacy mechanism for the SIP (Session Initiation Protocol)[RFC3261], based on techniques available at the time of its publication. This mechanism relies on the use of a separate privacy service to remove privacy-sensitive information from SIP messages sent by a user agent before forwarding those messages to the final destination. Since then, numerous SIP extensions have been proposed and standardized. Some of those enable a user agent to withhold its user's identity and related information without the need for privacy services, which was not possible when RFC 3323 was defined.

The purpose of this document is not to obsolete RFC 3323, but to enhance overall privacy mechanism in SIP by allowing user agent to take control of its privacy, rather than being completely dependent on external privacy service.

The UA-driven privacy mechanism defined in this document will not eliminate the need for RFC 3323 usage defined in [RFC3325] which instructs privacy service to delete P-Asserted-Identity header. In order to delete a P-Asserted-Identity header, a user agent needs to set Privacy:id even when the user agent is utilizing this specification.

This document defines a best current practice in which a user agent controls all the privacy functions on its own utilizing SIP extensions such as GRUU (Globally Routable User Agent URIs) [I-D.ietf-sip-gruu] and TURN (Traversal Using Relays around NAT) [I-D.ietf-behave-turn].

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 [RFC2119].

privacy-sensitive information:
The information that identifies a user who sends the SIP message, as well as the supplementary information that can be used to guess the user's identity.

3. Concept of Privacy

The concept of privacy in this document is the act of concealing identity of a user and supplementary information. The protection of network privacy (e.g., topology hiding) is outside the scope of this

document.

Privacy-sensitive information includes display-name and URI (Uniform Resource Identifier) in a From header field that can reveal the user's name and affiliation (e.g., company name), and IP addresses or host names in a Contact header field, a Via header field, a Call-ID header field, or an SDP (Session Description Protocol) [RFC4566] body that might reveal the location of a user agent.

4. Treatment of Privacy-Sensitive Information

Some fields of SIP messages that potentially contain privacy-sensitive information do not interfere with establishing dialog and can be omitted without any side effects. Other fields are essential for establishing dialog and need to have its value replaced to anonymized values in order to obscure privacy-sensitive information. Of the privacy-sensitive information listed in section 3, URIs, host names, and IP addresses in Contact, Via, and SDP must be functional (i.e., suitable for purpose) even when they are anonymized.

With the use of GRUU [I-D.ietf-sip-gruu] and TURN [I-D.ietf-behave-turn], a user agent can obtain URIs and IP addresses for media and signaling that are functional yet anonymous, and do not identify either the user agent or the user. Instructions on how to obtain a functional anonymous URI and IP address are given in Section 4.1 and 4.2, respectively.

Host names should be concealed because the user's identity may be

guessed from them, but they are not always regarded as critical privacy-sensitive information.

In addition, a user agent should be careful not to include any information that identifies the user in optional SIP header fields such as Subject and User-Agent.

4.1. Obtaining a Functional Anonymous URI Using the GRUU Mechanism

A user agent wanting to obtain a functional anonymous URI MUST support and utilize the GRUU mechanism unless it is able to obtain a functional anonymous URI through other means outside the scope of this document. By sending a REGISTER request requesting GRUU, the user agent can obtain an anonymous URI, which can later be used for the Contact header field.

The detailed process on how a user agent obtains a GRUU is described in [I-D.ietf-sip-gruu].

If the Registrar supports the GRUU and returns a REGISTER response, the user agent SHOULD search within the REGISTER response for a "temp-gruu" URI parameter, which provides the desired privacy property. If the "temp-gruu" URI parameter and value is present within the REGISTER response, the user agent SHOULD use the value of the "temp-gruu" as an anonymous URI representing the user agent. This URI SHOULD be used for the Contact header in subsequent requests and responses.

If there is no "temp-gruu" URI parameter in the 200 response to the REGISTER request, a user agent SHOULD NOT proceed with its anonymization process, unless something equivalent to "temp-gruu" is provided through some administrative means.

It is RECOMMENDED that user agent consult the user before sending a request without a functional anonymous URI when privacy is requested from the user.

Due to the nature of how GRUU works, the domain name is always revealed when GRUU is used. If revealing the domain name in Contact header field is a concern, usage of a third-party GRUU server to obtain a temp-gruu that is irrelevant to users' domain SHOULD be

considered. Refer to the Security Considerations section for details.

4.2. Obtaining a Functional Anonymous IP Address Using the TURN Mechanism

A user agent that is not provided with a functional anonymous IP address through some administrative means, MUST obtain a relayed address if anonymity is desired (IP address of the media relay) for use in SDP and in Via header. Such IP address is to be derived from a STUN relay server through TURN mechanism, which allows a STUN server to act as a media relay.

Anonymous IP addresses are needed for two purposes. The first is for use in the Via header field of a SIP request. By obtaining an IP address from a STUN relay server, using that address in the Via header field of the SIP request, and sending the SIP request to the STUN relay server, the IP address of the user agent will not be revealed beyond the relay server.

The second is for use in SDP as an address for receiving media. By obtaining an IP address from a STUN relay server and using that address in SDP, media will be received via the relay server. Also media can be sent via the relay server. In this way, neither SDP nor media packets reveal the IP address of the user agent.

It is assumed that a user agent is either manually or automatically configured through means such as the configuration framework [I-D.ietf-sipping-config-framework] with the address of one or more STUN (Session Traversal Utilities for NAT) [I-D.ietf-behave-turn] relay servers to obtain anonymous IP address.

5. User Agent Behavior

This section describes how to generate an anonymous SIP message at a user agent.

A user agent fully compliant with this document MUST obscure or conceal all the critical UA-inserted privacy-sensitive information in SIP requests and responses as shown in Section 5.1 when user privacy

is requested. In addition, the user agent SHOULD conceal the non-critical privacy-sensitive information as shown in Section 5.2.

Furthermore, when a user agent uses relay server to conceal its identity, the user agent MUST send requests to the relay server to ensure request and response bypass the same signaling path.

5.1. Critical Privacy-Sensitive Information

5.1.1. Contact Header Field

Without privacy considerations, this field contains a URI used to reach the user agent for mid-dialog requests and possibly out-of-dialog requests, such as REFER [RFC3515]. The Contact header field can also contain a display-name. Since the Contact header field is used for routing further requests to the user agent, it must include a functional URI even when it is anonymized.

A user agent generating an anonymous SIP message supporting this specification MUST anonymize a Contact header using an anonymous URI ("temp-gruu") obtained through the GRUU mechanism or an anonymous URI containing an IP address obtained through the TURN mechanism, unless an equivalent functional anonymous URI is provided by some other means.

Refer to Section 4.1 for details on how to obtain an anonymous URI through GRUU, and refer to Section 4.2 for details on how to obtain an IP address through TURN.

A display-name in a Contact header MUST be omitted or "Anonymous".

5.1.2. From Header Field

Without privacy considerations, this field contains the identity of the user, such as display-name and URI.

RFCs 3261 and 3323 recommend to set "sip:anonymous@anonymous.invalid" as a SIP URI in a From header field when user privacy is requested. This raises an issue when the SIP-Identity mechanism [RFC4474] is

applied to the message, because SIP-Identity requires an actual domain name in the From header field.

A user agent generating an anonymous SIP message supporting this specification MUST anonymize the From header field in one of the two ways described below.

Option 1:

A user agent anonymizes a From header field using an anonymous display-name and an anonymous URI following the procedure noted in section 4.1.1.3 of RFC 3323.

The example form of the From header of option 1 is as follows:

```
From: "Anonymous" <sip:anonymous@anonymous.invalid>;tag=1928301774
```

Option 2:

A user agent anonymizes a From header field using an anonymous display-name and an anonymous URI with user's valid domain name instead of "anonymous.invalid".

The example form of the From header of option 2 is as follows:

```
From: "Anonymous" <sip:anonymous@atlanta.com>;tag=1928301774
```

A user agent SHOULD go with option 1 to conceal its domain name in From header field. However, the SIP-Identity will fail with the From header of option 1 because the SIP-Identity mechanism uses authentication based on the domain name.

If a user agent is aware that SIP-Identity mechanism will be applied to the request, it is RECOMMENDED to go with option 2. However, the user's domain name will be revealed from the From header field of option 2.

If user wants both anonymity and strong identity, use a third party anonymization service which issues AoR for the use in From address which also provides SIP-Identity.

5.1.3. Via Header Field

Without privacy considerations, the bottommost Via header field added by a user agent contains the IP address and port or hostname that are used to reach the user agent for responses.

A user agent generating an anonymous SIP request supporting this specification MUST anonymize the IP address in the Via header field using an anonymous IP address obtained through the TURN mechanism, unless an equivalent functional anonymous IP address is provided by some other means.

Refer to Section 4.2 for details on how to obtain an IP address through TURN.

Via header SHOULD NOT include a host name.

5.1.4. IP Addresses in SDP

A user agent generating an anonymous SIP message supporting this specification MUST anonymize IP addresses in SDP, if present, using an anonymous IP address obtained through the TURN mechanism, unless an equivalent functional anonymous IP address is provided by some other means.

Refer to Section 4.2 for details on how to obtain an IP address through TURN.

5.2. Non-Critical Privacy-Sensitive Information

5.2.1. Host Names in Other SIP Header Fields

A user agent generating an anonymous SIP message supporting this specification SHOULD conceal host names in any SIP header fields, such as Call-ID and Warning header fields, if considered privacy-sensitive.

5.2.2. Optional SIP Header Fields

Other optional SIP header fields (such as Call-Info, In-Reply-To, Organization, Referred-By, Reply-To, Server, Subject, User-Agent, and Warning) can contain privacy-sensitive information.

A user agent generating an anonymous SIP message supporting this specification SHOULD NOT include any information that identifies the user in such optional header fields.

6. Security Considerations

This specification uses GRUU and TURN and inherits any security considerations described in these drafts.

Furthermore, if the provider of the caller intending to obscure its identity consists of a small number of people (e.g. small enterprise, SOHO), the domain name alone can reveal the identity of the caller when this specification is used.

Same can be true when the provider is large, but the receiver of the call only knows few people from the source of call.

There are mainly two places in the message, From header and Contact header, where domain name must be functional.

The domain name in From header can be obscured as described in section 5.1.2, on contrary the Contact header needs to contain a valid domain name at all time to function properly.

It is probably important to note that generally a device will not show the contact address to the receiver, but this does not mean that one can not find the domain name in a message. In fact as long as this specification is used to obscure identity, the message will always contain a valid domain name as it inherits key characteristics of GRUU.

If one wants to assure anonymization, it is recommended for the user to seek and rely on third party anonymization service.

A third party anonymization service provides registrar and TURN service which has no affiliation with the caller's provider, allowing caller to completely withhold its identity.

7. IANA Considerations

This document requires no action by IANA.

8. References

8.1. Normative References

[I-D.ietf-behave-turn]
Rosenberg, J., Mahy, R., and P. Matthews, "Traversal Using

draft-ietf-behave-turn-11 (work in progress),
October 2008.

[I-D.ietf-sip-gruu]
Rosenberg, J., "Obtaining and Using Globally Routable User Agent (UA) URIs (GRUU) in the Session Initiation Protocol (SIP)", draft-ietf-sip-gruu-15 (work in progress), October 2007.

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

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

[RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.

8.2. Informative References

[I-D.ietf-sipping-config-framework]
Channabasappa, S., "A Framework for Session Initiation Protocol User Agent Profile Delivery", draft-ietf-sipping-config-framework-15 (work in progress), February 2008.

[RFC3323] Peterson, J., "A Privacy Mechanism for the Session Initiation Protocol (SIP)", RFC 3323, November 2002.

[RFC3325] Jennings, C., Peterson, J., and M. Watson, "Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks", RFC 3325, November 2002.

[RFC3515] Sparks, R., "The Session Initiation Protocol (SIP) Refer Method", RFC 3515, April 2003.

[RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, August 2006.

Authors' Addresses

Mayumi Munakata
NTT Corporation

Email: munakata.mayumi@lab.ntt.co.jp

Shida Schubert
NTT Corporation

Email: shida@ntt-at.com

Takumi Ohba
NTT Corporation
9-11, Midori-cho 3-Chome
Musashino-shi, Tokyo 180-8585
Japan

Phone: +81 422 59 7748
Email: ohba.takumi@lab.ntt.co.jp
URI: http://www.ntt.co.jp

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.