              The SIP P-Served-User Private-Header (P-Header)
        for the 3GPP IP Multimedia (IM) Core Network (CN) Subsystem

Status of This Memo

   This memo provides information for the Internet community. It does
   not specify an Internet standard of any kind. Distribution of this
   memo is unlimited.

Copyright Notice

Abstract

   This document specifies the SIP P-Served-User P-header. This header
   field addresses an issue that was found in the 3rd Generation
   Partnership Project (3GPP) IMS (IP Multimedia Subsystem) between an
   S-CSCF (Serving Call Session Control Function) and an AS (Application
   Server) on the ISC (IMS Service Control) interface. This header
   field conveys the identity of the served user and the session case
   that applies to this particular communication session and application
   invocation.

Table of Contents

1. Introduction

   The 3rd Generation Partnership Project (3GPP) IMS (IP Multimedia
   Subsystem) uses SIP (RFC 3261 [2]) as its main signaling protocol.
   (For more information on the IMS, a detailed description can be found
   in 3GPP TS 23.228 [9] and 3GPP TS 24.229 [11].) 3GPP has identified
   issues with the linking in of a SIP application server that are most
   appropriately resolved by defining a new SIP P-header, according to
   the procedures in RFC 3427 [5].

   The remainder of this document is organized as follows. Section 4
   outlines the problem by using particular service scenarios, and
   Section 5 discusses the requirements derived from these scenarios.
   Section 6 defines the P-Served-User header field, which meets those
   requirements, Section 7 specifies the proxy behavior for the new
   header field, and Section 8 discusses the applicability and scope of
   this new header field. Section 9 registers the P-Served-User header
   field with the IANA, and Section 10 discusses the security properties
   of the environment where this header field is intended to be used.

2. Conventions

   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 BCP 14, RFC 2119 [1].

3. Definitions

3.1. Identity, Network Asserted Identity, Trust Domain, and Spec(T)

   The terms Identity, Network Asserted Identity, Trust Domain, and
   Spec(T) in this document are specified in RFC 3324 [3].

3.2. Served User

   The served user to a proxy or AS (Application Server) is the user
   whose service profile is accessed by that proxy or AS when an initial
   request is received that is originated by, originated on behalf of,
   or terminated to that user. This profile in turn provides some
   useful information (preferences or permissions) for processing at a
   proxy and, potentially, at an AS.

4. Scenarios

4.1. General

   In the 3GPP IMS (IP Multimedia Subsystem), the S-CSCF (Serving CSCF)
   is a SIP proxy that serves as a registrar and handles originating and
   terminating session states for users allocated to it. This means
   that any call that is originated by a specific user or any call that
   is terminated to that specific user will pass through the S-CSCF that
   is allocated to that user.

   At the moment that an S-CSCF is allocated for a specific user, a user
   profile is downloaded to the S-CSCF from the HSS (Home Subscriber
   Server) over the Cx interface, see 3GPP TS 29.228 [12]. This user
   profile tells the S-CSCF whether the user is allowed to originate or
   terminate calls or whether an AS needs to be linked in over the ISC
   interface. The user profile information that determines whether a
   particular initial request needs to be sent to a particular AS is
   called the initial Filter Criteria (iFC), see for example 3GPP TS
   23.218 [8].

   For an S-CSCF to be able to meet its responsibilities, it needs to
   determine on which user's behalf it is performing its tasks and which
   session case is applicable for the particular request. (For a
   definition of session case, see 3GPP TS 29.228 [12]). The session
   case distinguishes the originating and terminating call cases and
   determines whether or not the particular user is registered.

   When the S-CSCF determines that for an incoming initial request the
   originating call case applies, it determines the served user by
   looking at the P-Asserted-Identity header field (RFC 3325 [4]), which
   carries the network asserted identity of the originating user. When,
   after processing the iFC for this initial request, the S-CSCF decides
   to forward the request to an AS, the AS has to go through a similar
   process of determining the session case and the served user. Since
   it should come to the same conclusion that this is an originating
   session case, it also has to look at the P-Asserted-Identity header
   field to determine the served user.

   When the S-CSCF determines that for an incoming initial request the
   terminating call case applies, it determines the served user by
   looking at the Request-URI (RFC 3261 [2]), which carries the identity
   of the intended terminating user. When, after processing the iFC for
   this initial request, the S-CSCF decides to forward the request to an
   AS, the AS has to go through a similar process of determining the
   session case and the served user. Since it should come to the same
   conclusion that this is a terminating session case, it also has to
   look at the Request-URI to determine the served user.

In the originating case, it can be observed that while the
P-Asserted-Identity header field just represents the originating user
when it enters the S-CSCF, it is overloaded with another meaning when
it is sent to an AS over the ISC interface. This other meaning is
that it serves as a representation of the served user.

In the terminating case, a similar overloading happens to the
Request-URI; while it first only represented the identity of the
intended terminating user, it is overloaded with another meaning when
it is sent to an AS over the ISC interface. This other meaning is
that it serves as a representation of the served user.

In basic call scenarios, this does not show up as a problem, but once
more complicated service scenarios (notably forwarding services) need
to be realized, it poses severe limitations. Such scenarios are
brought forward in the following subsections.

4.2. Diversion: Continue on Terminating Leg, but Finish Subsequent
      Terminating iFC First

   Imagine a service scenario where a user B has a terminating service
   that diverts the call to a different destination but is required to
   still execute subsequent terminating services for the same user.
   This means that this particular user has multiple iFC configured that
   are applicable for an incoming initial request. When the S-CSCF
   receives an initial INVITE request, it analyzes the request and
   determines that the session case is for a terminating registered
   user, then it determines the served user to be user B by looking at
   the Request-URI.

   Now the S-CSCF starts the iFC processing. The first iFC that matches
   the INVITE request causes the INVITE to be forwarded over the ISC
   interface to an AS that hosts user B's diversion service by adding
   the AS and S-CSCF's own hostnames to the Route header. The S-CSCF
   adds an Original Dialog Identifier (ODI) to the S-CSCF's own hostname
   on the Route header. This allows the S-CSCF to correlate an INVITE
   coming from an AS over the ISC interface to the existing session that
   forwarded the INVITE to the AS in the first place.

   When the AS receives the initial INVITE request, it analyzes the
   request and determines that the session case is for a terminating
   registered user, then it determines the served user to be user B by
   looking at the Request-URI. Based on some criteria, the diversion
   service concludes that the request needs to be diverted to another
   user or application C. It does this by changing the Request-URI to
   C. Optionally, it records the Request-URI history by using the
   History- Info header field (RFC 4244 [7]). Then the AS removes

itself from the Route header and routes the INVITE request back to
the S-CSCF by using the topmost Route header field.

When the S-CSCF receives the INVITE over the ISC interface, it can
see that the Route header contains its own hostname and an ODI that
correlates to an existing terminating session for user B. This can
be used by the S-CSCF to analyze whether there are still unexecuted
iFC. (Note that the current behavior of the S-CSCF on receiving an
INVITE with a changed Request-URI is to terminate the iFC processing
and to route the request based on the new Request-URI value.)

The process repeats itself. The INVITE is forwarded to the AS that
is associated with this particular iFC. When the AS receives the
initial INVITE request, it analyzes the request and determines that
the session case is for a terminating registered user, then it
determines the served user to be user C by looking at the Request-
URI. This is clearly wrong, as the user being served is still user
B.

This scenario clearly shows the problem that occurs when the Request-
URI is overloaded with the meanings "intended target identity" and
"served user" with the operation as described in Section 4.1. And it
shows that this use case can not be realized without introducing a
mechanism that conveys information about the served user from the
S-CSCF to the AS. Use of the History-Info element does not solve
this problem as it does not tell the AS which user is being served;
it just presents a history of diversions that might not be even
caused by the systems serving this particular user. A more detailed
analysis on why the History-Info header field can't be used is
provided in Appendix A.

4.3. Diversion: Create New Originating Leg and Provide Originating iFC
      Processing

   Imagine a service scenario where a user B has a terminating service
   that diverts the call to a different destination. It is required
   that a forwarded call leg is handled as an originating call leg and
   that originating services for user B are executed. This means that
   this particular user has one or more iFC configured that are
   applicable for an outgoing initial request.

   When the S-CSCF receives an initial INVITE request, it analyzes the
   request and determines that the session case is for a terminating
   registered user, then it determines the served user to be user B by
   looking at the Request-URI.

Now the S-CSCF starts the iFC processing. The first iFC that matches
the INVITE request causes the INVITE to be forwarded over the ISC
interface to an AS that hosts user B's diversion service by adding
the AS and S-CSCF's own hostnames to the Route header. The S-CSCF
adds an Original Dialog Identifier (ODI) to the S-CSCF's own hostname
on the Route header. This allows the S-CSCF to correlate an INVITE
coming from an AS over the ISC interface to the existing session that
forwarded the INVITE to the AS in the first place.

When the AS receives the initial INVITE request, it analyzes the
request and determines that the session case is for a terminating
registered user, then it determines the served user to be user B by
looking at the Request-URI. Based on some criteria, the diversion
service concludes that the request needs to be diverted to another
user or application C. It does this by changing the Request-URI to
C. Optionally, it records the Request-URI history by using the
History-Info header field (RFC 4244 [7]). Then the AS removes itself
from the Route header. To make sure that the request is handled as a
new originating call on behalf of user B, the AS adds the "orig"
parameter to the topmost route header. Then it routes the INVITE
request back to the S-CSCF by using this topmost Route header field.

When the S-CSCF receives the INVITE over the ISC interface, it can
see that the topmost Route header contains its own hostname and an
"orig" parameter. Because the topmost Route header contains the
"orig" parameter, the S-CSCF concludes that the INVITE should be
handled as if a call is originated by the served user. The served
user is determined from the P-Asserted-Identity header to be user A.
 This is clearly wrong, as the user being served is and should be
user B.

For the sake of discussion, let's assume that the S-CSCF can
determine that the served user is user B. Then the procedure would
continue as follows: The S-CSCF starts the originating iFC
processing, the first iFC that matches the INVITE request causes the
INVITE to be forwarded over the ISC interface to an AS that hosts an
originating service of user B by adding the AS and S-CSCF's own
hostnames to the Route header. The S-CSCF adds an Original Dialog
Identifier (ODI) to the S-CSCF's own hostname on the Route header.

The INVITE is forwarded to the AS that is associated with this
particular iFC. When the AS receives the initial INVITE request, it
analyzes the request and determines that the session case is for an
originating registered user, then it determines the served user to be
user A by looking at the P-Asserted-Identity. This is clearly wrong,
as the user being served is and should be user B.

   This scenario clearly shows the problem that occurs when the
   P-Asserted-Identity is overloaded with the meanings "call originator"
   and "served user" with the operation as described in Section 4.1.
   And it shows that this use case can not be realized without
   introducing a mechanism that conveys information about the served
   user from the S-CSCF to the AS and from the AS to the S-CSCF. Use of
   the History-Info element does not solve this problem as it does not
   tell the AS which user is being served, but just presents a history
   of diversions that might not be even caused by the systems serving
   this particular user. A more detailed analysis on why the History-
   Info header field can't be used is provided in Appendix A.

4.4. Call Out of the Blue: on Behalf of User B, but Service Profile of
      Service Identity C

   There are services that need to be able to initiate a call, whereby
   the call appears to be coming from a user B but the service profile
   on behalf of service identity C needs to be executed in the S-CSCF.

   When a call needs to appear as coming from user B, that means that
   the P-Asserted-Identity needs to contain B's identity. This is
   because the Originating Identity Presentation (OIP) service as
   defined in 3GPP TS 24.173 [10] uses the P-Asserted-Identity to
   present the call originator. This makes sense because that is the
   main meaning expressed by the P-Asserted-Identity header field.

   It is clear that no INVITE request can be constructed currently that
   would achieve both requirements expressed in the first paragraph,
   because the P-Asserted-Identity is overloaded with two meanings on
   the ISC interface. When the S-CSCF will receive this request, it
   will determine that the served user is user B, which is not what we
   want to achieve.

5. Requirements

   This section lists the requirements derived from the previous
   scenarios:

   1. To be able to offer real-world application services, it is
       required that the identity of the served user can be conveyed on
       the ISC interface (see 3GPP TS 23.218 [8]).

   2. To be able to offer appropriate services to the served user, it
       is required that in addition to the served user identity the
       session case is conveyed.

6. P-Served-User Header Field Definition

   This document defines the SIP P-Served-User P-header. This header
   field can be added to initial requests for a dialog or standalone
   requests, which are routed between nodes in a Trust Domain for
   P-Served-User. The P-Served-User P-header contains an identity of
   the user that represents the served user. The "sescase" parameter
   may be used to convey whether the initial request is originated by or
   destined for the served user. The "regstate" parameter may be used
   to indicate whether the initial request is for a registered or
   unregistered user.

   The augmented Backus-Naur Form (BNF) (RFC 5234 [6]) syntax of the
   P-Served-User header field is as follows:

   P-Served-User   = "P-Served-User" HCOLON PServedUser-value
                                 *(SEMI served-user-param)
   served-user-param   = sessioncase-param
                                 / registration-state-param
                                 / generic-param
   PServedUser-value   = name-addr / addr-spec
   sessioncase-param   = "sescase" EQUAL "orig" / "term"
   registration-state-param = "regstate" EQUAL "unreg" / "reg"

   EQUAL, HCOLON, SEMI, name-addr, addr-spec, and generic-param are
   defined in RFC 3261 [2].

   The following is an example of a P-Served-User header field:

   P-Served-User: <sip:user@example.com>; sescase=orig; regstate=reg

7. Proxy Behavior

7.1. Generating the P-Served-User Header

   Proxies that support the header MUST only insert the header in
   initial requests for a dialog or in standalone requests when the
   following conditions hold:

   o The proxy has the capability to determine the served user for the
      current request.

   o The next hop is part of the same Trust Domain for P-Served-User.

   When the above conditions do not hold, the proxy MUST NOT insert the
   header.

7.2. Consuming the P-Served-User Header

   A proxy that supports the header MUST, upon receiving from a trusted
   node the P-Served-User header in initial requests for a dialog or in
   standalone requests, take the value of the P-Served-User header to
   represent the served user in operations that require such
   information.

   A proxy that supports the header MUST remove the header from requests
   or responses when the header was received from a node outside the
   Trust Domain for P-Served-User before further forwarding the message.

   A proxy that supports the header MUST remove the header from requests
   or responses when the next hop is a node outside the Trust Domain for
   P-Served-User before further forwarding the message.

8. Applicability

   According to RFC 3427 [5], P-headers have a limited applicability.
   Specifications of P-headers, such as this RFC, need to clearly
   document the useful scope of the proposal and explain its limitations
   and why it is not suitable for the general use of SIP on the
   Internet.

   The use of the P-Served-User header field extensions is only
   applicable inside a Trust Domain for served user. Nodes in such a
   Trust Domain explicitly trust each other to convey the served user
   and to be responsible for withholding that information outside of the
   Trust Domain. The means by which the network determines the served
   user and the policies that are executed for a specific served user is
   outside the scope of this document.

   The served user information lacks an indication of who or what
   specifically determined the served user, and so it must be assumed
   that the Trust Domain determined the served user. Therefore, the
   information is only meaningful when securely received from a node
   known to be a member of the Trust Domain.

   Because the served user typically only has validity in one
   administrative domain, it is in general not suitable for inter-domain
   use or use in the Internet at large.

   Despite these limitations, there are sufficiently useful specialized
   deployments that meet the assumptions described above, and that can
   accept the limitations that result, to warrant informational
   publication of this mechanism. An example deployment would be a
   closed network like 3GPP IMS.

9. IANA Considerations

   This document defines a new SIP header field: P-Served-User. This
   header field has been registered by the IANA in the SIP Parameters
   registry under the Header Fields subregistry.

10. Security Considerations

   The P-Served-User header field defined in this document is to be used
   in an environment where elements are trusted and where attackers are
   not supposed to have access to the protocol messages between those
   elements. Traffic protection between network elements is sometimes
   achieved by using IPsec and sometimes by physically protecting the
   network. In any case, the environment where the P-Served-User header
   field will be used ensures the integrity and the confidentiality of
   the contents of this header field.

   The Spec(T) that defines the Trust Domain for P-Served-User MUST
   require that member nodes understand the P-Served-User header
   extension.

   There is a security risk if a P-Served-User header field is allowed
   to propagate out of the Trust Domain where it was generated. In that
   case, user-sensitive information would be revealed. To prevent such
   a breach from happening, proxies MUST NOT insert the header when
   forwarding requests to a hop that is located outside the Trust
   Domain. When forwarding the request to a node in the Trust Domain,
   proxies MUST NOT insert the header unless they have sufficient
   knowledge that the route set includes another proxy in the Trust
   Domain that understands the header, such as the home proxy. There is
   no automatic mechanism to learn the support for this specification.
   Proxies MUST remove the header when forwarding requests to nodes that
   are not in the Trust Domain or when the proxy does not have knowledge
   of any other proxy included in the route set that will remove it
   before it is routed to any node that is not in the Trust Domain.

11. Acknowledgments

   Alf Heidermark, Hubert Przybysz, and Erik Rolin for the discussion
   that led to the solution written down in this document. Spencer
   Dawkins for performing the expert review. Jon Peterson for
   performing the AD review. Gonzalo Camarillo, Paul Kyzivat, Nils
   Haenstroem, Arunachalam Venkatraman, Mikael Forsberg, Miguel Garcia,
   Jozsef Varga, Keith Drage, Tim Polk, and Cullen Jennings for
   providing improvements. Francis Dupont for performing the general
   area review. Sandy Murphy for performing the security area review.

12. References

12.1. Normative References

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

   [2]   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.

   [3]   Watson, M., "Short Term Requirements for Network Asserted
           Identity", RFC 3324, November 2002.

   [4]   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.

   [5]   Mankin, A., Bradner, S., Mahy, R., Willis, D., Ott, J., and B.
           Rosen, "Change Process for the Session Initiation Protocol
           (SIP)", BCP 67, RFC 3427, December 2002.

   [6]   Crocker, D. and P. Overell, "Augmented BNF for Syntax
           Specifications: ABNF", STD 68, RFC 5234, January 2008.

12.2. Informative References

   [7]   Barnes, M., "An Extension to the Session Initiation Protocol
           (SIP) for Request History Information", RFC 4244,
           November 2005.

   [8]   3GPP, "IP Multimedia (IM) session handling; IM call model;
           Stage 2", 3GPP TS 23.218 V7.

   [9]   3GPP, "IP Multimedia Subsystem (IMS); Stage 2", 3GPP TS 23.228
           V7.

   [10]   3GPP, "IMS multimedia telephony communication service and
           supplementary services; Stage 3", 3GPP TS 24.173 V7.

   [11]   3GPP, "Internet Protocol (IP) multimedia call control protocol
           based on Session Initiation Protocol (SIP) and Session
           Description Protocol (SDP); Stage 3", 3GPP TS 24.229 V7.

   [12]   3GPP, "IP Multimedia (IM) Subsystem Cx and Dx interfaces;
           Signalling flows and message contents", 3GPP TS 29.228 V7.

Appendix A. Why the History-Info Header Is Not Suitable to Convey the
             Served User Information on the ISC Interface

A.1. Semantics

   The History-Info (as specified in RFC 4244 [7]) holds a record of
   subsequent Request-URI values that are put on an initial request
   during its processing in the network.

   If it would be possible at all to use the History-Info header for the
   purpose of communicating the served user, then again the same
   overloading would occur as the one that we are trying to get rid of
   (Section 4.2). In this case, we overload the particular History-Info
   header field's hi-entry with the meaning "historic target identity"
   and "served user".

   Another reason that the History-Info header can not solve the
   requirements as expressed in this document is that, in originating
   session case scenarios, the served user is currently determined from
   the P-Asserted-Identity, as that header field contains the asserted
   originating user's identity. The History-Info header, being a record
   of Request-URIs, can never be a solution for this case.

   Looking at the call-out-of-the-blue scenario (Section 4.4), it is
   impossible to construct a History-Info header for an INVITE request
   on behalf of user C that appears to come from user B and targets user
   D that would express the served user C without violating the original
   semantics of the History-Info header according to (RFC 4244 [7]).

A.2. Additional Observations

   The purpose of the History-Info header is a header that has an end-
   to-end application. For the purpose of informing an AS on the ISC
   interface, this is overkill.

   At the moment that the AS receives an initial INVITE over the ISC
   interface, this INVITE may have passed a vast number of proxies that
   may or may not have added history information. On top of that, the
   request may have traversed several AS instances for the same served
   user. In case several subsequent iFC are active, all these AS
   instances may perform a forwarding. This means that it is not
   possible to define an algorithm that points out which hi-entry of a
   History-Info header should represent the served user. In other
   words, a History-Info header field with n entries expresses a branch
   of depth n. Any or none of these elements could be the served user
   identity.

   The History-Info header does not comply with the second requirement
   as expressed in Section 5, as it does not have a means to express the
   session case in a natural way.

A.3. Conclusion

   Each observation in the previous subsections, alone, is enough to
   disregard the History-Info header as an information element that is
   suitable for transporting the served user information over the ISC
   interface.

   Note that this does not prohibit the use of the P-Served-User header
   and the History-Info header in the same request. In fact that will
   be a quite likely scenario for network-based diversion services like,
   for example, the Communication Diversion service as specified in
   (3GPP TS 24.173 [10]).

Author's Address

   Hans Erik van Elburg
   Ericsson Telecommunicatie B.V.
   Ericssonstraat 2
   Rijen 5121 ML
   Netherlands

   EMail: HansErik.van.Elburg@ericsson.com