

Dispatch	K.O. Ono
Internet-Draft	H.S. Schulzrinne
Intended status: Standards Track	Columbia University
Expires: April 23, 2012	October 21, 2011

Referencing and Validating User Attributes
draft-ono-dispatch-attribute-validation-00.txt

Abstract

This document describes a mechanism for referencing and validating user attributes in SIP communication. User attributes, such as an organizational affiliation and role, are helpful for the recipients of a communication request to decide whether or not to grant the sender access to the recipient's resources, especially when the sender identity is unknown to the recipients. This mechanism allows the sender to claim her attributes to recipients using an attribute reference identifier without needing to prove the sender identity. This document defines a new SIP "Sender-References" header field to convey one or more attribute reference identifiers. This mechanism satisfies all the requirements for trait-based authorization defined in RFC 4484, except that it provides only one assertion scheme.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet- Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 23, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

[Table of Contents](#)

- *1. [Introduction](#)
- *2. [Terminology](#)
- *3. [Architecture](#)
 - *3.1. [Assumed Trust Relationships among AVS, Caller, and Callee](#)
 - *3.2. [ARIDs are Loosely Associated with the Owner's Identity in SIP](#)
- *4. [Requirements](#)
 - *4.1. [Differences between Our Requirements and the Requirements for Trait-Based Authorization](#)
- *5. [Procedures](#)
 - *5.1. [Generating an ARID](#)
 - *5.2. [Obtaining an ARID](#)
 - *5.3. [Sending an ARID in a Communication Request](#)
 - *5.4. [Validating an ARID to Retrieve User Attributes](#)
- *6. [Sender-References Header Field](#)
- *7. [Relationship to Existing Mechanisms](#)
- *8. [Security Considerations](#)
 - *8.1. [Man in the Middle Attacks](#)
 - *8.2. [Replay Attacks Using a Received ARID](#)
 - *8.3. [Denial of Service Attacks on the AVS](#)
 - *8.4. [Phishing Attacks on the AVS](#)
- *9. [IANA Considerations](#)
- *10. [References](#)
 - *10.1. [Normative References](#)
 - *10.2. [Informative References](#)
- *[Authors' Addresses](#)

1. Introduction

Ascertaining a person's attributes is often useful to determine the trustworthiness of the person when two people first meet each other. These user attributes include, for example, an organizational affiliation, a role in a professional society, age, holding certificates or licenses, and being a customer of a bank, an employee, or a student. If user attributes are available with a communication request, these attributes can help the recipient determine how to handle the communication request by estimating whether the communication is important enough to be established.

A caller identifier (ID) authenticated by the SIP Identity mechanism [\[RFC4474\]](#), when used alone, can be a helpful user attribute, but only in limited cases. Only if a caller ID is in a SIP-URI [\[RFC3261\]](#) and is authenticated by the domain of a trusted organization can the caller ID be perceived as evidence that the caller belongs to the trusted organization. However, if a caller ID in a SIP-URI belongs to an untrusted domain regarding user admission policy, such as a free voice over IP service provider, or if a caller ID does not contain any domain name, such as a tel-URI [\[RFC3966\]](#), the caller ID does not indicate the caller's trustworthiness to the callee who has never seen the caller ID before. Thus, even if a caller has multiple contact addresses, the caller needs to use a contact address issued by a trusted domain for authorization purposes. To offer a flexible choice of which contact address to use, our referencing mechanism introduces another piece of information, an attribute reference ID (ARID), that enables a caller to refer to her attributes without needing to rely on the caller ID. A caller can use multiple ARIDs if the caller wants to prove multiple attributes associated with different organizations. This referencing mechanism, unlike the caller ID, allows a caller to use multiple ARIDs to declare multiple user attributes in a single communication request. If an authenticated caller ID does not provide sufficient information, the callee can look up further user attributes through directory services. However, a reference integrity problem arises when a directory service does not allow queriers to look up user attributes by the user's contact address. Additionally, when a directory service allows queries by a user's contact address, but is offered by a third party, not the issuer of contact addresses, the authenticity of the information is unreliable. For example, DoctorFinder service offered by the American Medical Association provides information about certified medical doctors. When making a query, a querier cannot use the doctor's phone number, but needs to use doctor's common name, street address or specialty, which is available to the public. If a doctor makes a call (or sends an email message) that includes such query information and a reference to the DoctorFinder service, the callee (or the recipient) is not convinced of the certainty. To solve this reference integrity problem, our referencing mechanism allows an organization to generate a short-lived ARID upon a caller request. This ARID is effective only for a specific communication by limiting the lifetime and encoding

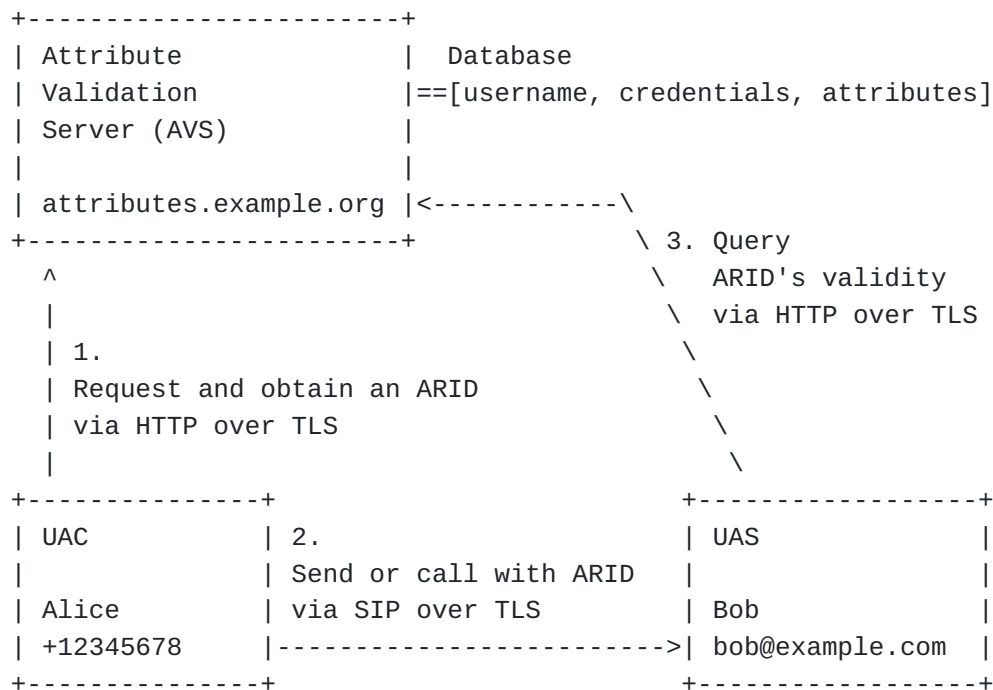
designated destinations, namely designated queriers. In addition, the ARID can be used only for retrieving the attributes that the caller selects to disclose to the specific queriers.

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 [\[RFC2119\]](#).

3. Architecture

[Figure 1](#) depicts an overview of the service architecture where an attribute validation server (AVS) operates to reference and validate user attributes for an organization. For each user, the AVS maintains the username and credentials to authenticate the user for remote access, in addition to other information such as a user number and role which the organization assigns, the users's common name, affiliation, street address, and electronic contact addresses. Note that the AVS stores a user's contact addresses, but it neither guarantees that the user owns the contact addresses nor can be reached by their addresses. We provide an example for illustration. Alice, a user of services provided by the organization, "example.org", is about to make a call to Bob at "bob@example.com". Alice first requests an ARID from the AVS using HTTP [\[RFC2616\]](#) over TLS [\[RFC5246\]](#). When sending the request, Alice authenticates the AVS using its X.509 Public Key Certificate (PKC) [\[RFC5280\]](#) which is delivered in the TLS handshake and is signed by a trusted Certificate Authority (CA). In turn, when generating an ARID for Alice, the AVS authenticates her using any credentials supported by the AVS, such as a password or a client's X.509 PKC. Upon successfully obtaining an ARID, Alice makes a call to Bob using SIP [\[RFC3261\]](#) over TLS. The SIP INVITE request includes the ARID. When Bob receives an ARID, he queries the validity of the ARID to the AVS using HTTP over TLS. Bob authenticates the AVS using its X.509 PKC in the same way that Alice does. Based on the query results, Bob determines whether or not to answer the call from Alice and adjusts his communication stance accordingly.



3.1. Assumed Trust Relationships among AVS, Caller, and Callee

We assume that the AVS and the caller, Alice, trust each other regarding the attribute validation service for an organization, "example.org." They share Alice's username and credentials for remote access, and her attributes. Alice trusts the AVS to properly maintain her attributes and to disclose the attributes she selects only to queriers whom she specifies. In turn, the AVS trusts Alice as a user in the organization and trusts her attributes which it knows first-hand, such as "Alice is an IEEE student member." However, the AVS does not know the authenticity of her attributes that are not issued by the organization, such as her common name, affiliation, and contact addresses.

We also assume that Bob knows "example.org" as the domain name of an organization that has a user admission policy he trusts, whether or not he belongs to the organization. Bob also trusts the AVS to properly perform its attribute validation service.

Alice finds Bob worth making a call and disclosing her attributes to establish a communication with him. In turn, Bob does not have sufficient information about Alice's trustworthiness based solely on her identity in a SIP communication request.

3.2. ARIDs are Loosely Associated with the Owner's Identity in SIP

An ARID generated upon Alice's request can be used only to retrieve her attributes, but the ARID is not tightly linked with her identity used in a SIP communication request, namely the caller ID in a call. When Bob receives an ARID in a SIP communication request where the message integrity is protected by TLS, the callee can perceive the ARID to be

associated with the caller ID. Bob can loosely link an ARID with the owner's identity only because of the fact that these two pieces of information are sent in the same message. Other than the presence of these two pieces of information in the same message, there is no linkage between the ARID and the caller ID. Bob does not need to provide Alice's caller ID to validate a received ARID. The user attributes Bob retrieves upon the success of the validation do not contain the owner's contact address. This loose linkage is a natural consequence of the general fact that user attributes and the user identifier in a communication are often issued separately by different organizations or services.

This loose linkage, however, makes it difficult for Bob to detect impersonation using a stolen ARID. Bob cannot detect this impersonation by providing the AVS with the owner's caller ID or by being presented the caller ID in user attributes. When issuing an ARID, the AVS cannot easily authenticate her caller ID since the caller ID is issued by a different administrative domain. Additionally, Bob cannot always authenticate the caller ID. The cases where no authentication of the caller ID is available include where a caller ID is in a SIP-URI issued by the domain which does not deploy the SIP Identity mechanism, where a caller ID is in a tel-URI which is sent without any other authentication mechanisms, such as a digital signature in S/MIME [\[RFC5751\]](#), and where a caller ID is anonymized. Thus, although tightening this linkage can protect from impersonation attacks, it makes the service deployment more difficult and limits the caller's choice of caller IDs.

To mitigate the vulnerability to impersonation attacks using a stolen ARID without tying an ARID to an authenticated caller ID, a countermeasure is devised for each vulnerable target. To prevent a man in the middle from eavesdropping on an ARID, all the connection links to convey an ARID need to be protected with TLS. To detect that an ARID was stolen from the owner, the recipients, or intermediaries, such as a SIP proxy server, an ARID can be used to retrieve user attributes only a limited number of times, for a limited time period, and by limited queriers. Yet even with these protections, this mechanism cannot prevent the owner of an ARID from giving her own ARID to others. To keep this mechanism simple, we do not include any additional mechanisms that discourage the owner from giving her own ARID. As a result, this mechanism allows the owner of an ARID to informally delegate her attributes to others without proving the chain of authorizing delegation. However, a legitimate recipient cannot impersonate Alice's attributes by forwarding a received ARID.

4. Requirements

This section first identifies the requirements of a mechanism for referencing and validating attributes, and then identifies differences between these requirements and the requirements for Trait-Based Authorization (TBA) for SIP [\[RFC4484\]](#).

Our requirements are:

We intentionally omit the following requirements:

[4.1. Differences between Our Requirements and the Requirements for Trait-Based Authorization](#)

Our requirements described above are similar to the TBA requirements for SIP, but two differences exist. First, we do not require support for optional assertion schemes other than an ARID defined in [Section 5](#) while the TBA includes the following requirement:

Our mechanism currently does not support other assertion schemes, such as SAML [\[SAML\]](#) or X.509 Attribute Certificates (AC) [\[RFC5755\]](#), as mentioned above. Such mechanisms that protect assertion integrity by signing using the issuer's private key requires that recipients verify the integrity using the issuer's public key in the application layer. The recipients also need to authenticate the issuer of an assertion. On the other hand, our mechanism relies on transport layer security, namely TLS, to protect message integrity and authenticate the issuer of an ARID. Although our mechanism does not separately protect the integrity of user attributes or the linkage between user attributes and their owner, our mechanism instead protects the integrity of a whole message including these attributes. As long as intermediaries such as an HTTP and SIP proxy servers can be trusted to properly transfer messages for this attribute referencing service, this security with TLS is simpler, and strong enough against message tampering and server impersonation.

The second difference is that our requirements include an additional requirement for protecting user's privacy described in REQ-8. Although an authorization service or AVS needs to limit designated queriers to the designated destinations of a SIP request, the authorization service has to know neither user's communication history nor plans containing routable contact addresses to do so even for a short term during the lifetime of an assertion or ARID. Our mechanism hashes contact addresses to prevent this unnecessary disclosure of the private information of a user.

[5. Procedures](#)

[Figure 2](#) illustrates message exchanges among a UAC, the UAS and the AVS for the following procedures:

Before explaining each procedure, we describe how the AVS typically generates an ARID.

Alice	AVS	Bob
UAC		UAS
F1. HTTP POST		
----->		
F2. 200 OK with ARID		
<-----		
F3. SIP INVITE with ARID		
----->		
	F4. HTTP GET with ARID	
	<-----	
	F5. 200 OK	
	----->	
F6. 200 OK		
<-----		
F7. ACK		
----->		

Note: SIP messages to/from SIP proxy servers are omitted since they are not affected by this mechanism.

5.1. Generating an ARID

An ARID is a string of URL [\[RFC3986\]](#) characters generated by an AVS upon a user's request. When a single AVS offers this attribute service for multiple organizations, a subdomain or a path in the URL of the AVS website is assigned to each organization as part of an ARID to meet the requirement REQ-9.

We show two examples how an AVS generates an ARID. Note that the AVS does not have to follow these generating mechanisms. The first example is to hash a string of characters by SHA1 [\[SHA1\]](#). The string of characters is a user number concatenated with the timestamp, a nonce, and hashed contact addresses of one or more desired queriers (REQ-4,8) as shown below. Hashed contact addresses of one or more desired queriers are sent from a user when the user requests an ARID, as described in [Section 5.2](#). The information other than these hashed contact addresses is stored or generated on the AVS.

Another example is to encrypt a string of characters with a symmetric key of the AVS using AES [\[AES\]](#). The string of characters is a user number concatenated with a disclosure mode, the expiry time, hashed contact addresses of desired queriers. The disclosure mode is determined what attributes a user discloses to desired queriers (REQ-3). The expiry time of an ARID needs to be shortly after the time

an ARID is generated, such as ten minutes later, to avoid replay attacks (REQ-7).

*An appropriate expiry time depends on the service type. For synchronous communication services, such as a voice or video call or real-time text chat, the lifetime needs to be short. For asynchronous services, such as instant messaging, or email communication, the lifetime needs to be longer, such as 24 hours.

Generating an ARID by encryption:

```
m = user number || disclosure mode
|| expiry time || salt|| hashed querier's contact address
```

If two queriers, querier_1 and querier_2, are specified,
m = user_id || disclosure_mode || expiry time || salt || hashed
querier_1's contact address;hashed querier_2's contact address

ARID = URL path/Encrypt(m)

When selecting a method for generating an ARID, by hash or encryption, they have the trade-off between the memory cost of storing ARIDs with related data and the computational cost of decrypting ARIDs. When generating an ARID by hash, the AVS needs to store the generated ARID with associated data including the expiry time, the nonce, the hashed contact addresses of desired queriers which the user sent, and the disclosure mode which the user specified. On the other hand, when generating an ARID by encryption, the AVS only needs to remember the salt for decryption, but not any generated ARIDs. Instead, it requires the computational cost of decryption.

5.2. Obtaining an ARID

To obtain an ARID which can be used for a communication with Bob, Alice first needs to connect to the AVS using a SIP UA which supports this mechanism. When connecting, the SIP UAC MUST authenticate the AVS using its X.509 PKC sent in the TLS handshake. In turn, the AVS MUST authenticate Alice using her username and credentials. For user authentication, HTTP Basic or Digest authentication [\[RFC2617\]](#), a client's PKC, or other mechanisms SHOULD be used. Upon successful user authentication, the SIP UAC MUST send the AVS an HTTP POST request with setting hashed Bob's contact address as a desired querier, and a disclosure mode in a message body, as shown in the following example. Each hashed contact address of a desired querier SHOULD be attached as a JSON [\[RFC4627\]](#) object, or MAY be in XML [\[XML\]](#). When a communication request has multiple destinations, such as a conference call, multiple "destination" fields SHOULD be included to contain multiple hashed contact addresses of the desired queriers.

F1. HTTP POST sent from Alice to AVS:

Hashing the contact address of a desired querier is to limit acceptable queriers without revealing communication history to the AVS (REQ-8).

The SIP UA supporting this mechanism MUST implement and use SHA1, and MAY support any other hash algorithms. To prevent re-identification based on hashed contact addresses collected on the AVS, the SIP UAC MUST generate a salt, which is a random string of characters, and concatenate it with a contact address as follows:

In the example above, the destination field, "2cf6a1eda3b5205005d25a7d5dcf13bb200fc26a", is generated by SHA1("dmvb1p03"||"sips:bob@example.com").

When the AVS successfully generates an ARID for Alice, the AVS responds to her with a 200 OK response including the ARID and its expiry time in the same data format used in the received HTTP request. The HTTP messages MUST be sent over TLS to protect message confidentiality and integrity. In the following example, the ARID is attached as a JSON object. The "arid" field consists of the URL of the website for the ARID validation, "https://attributes.example.org/", and the ARID, "17750c5cbac9979171991d505d2e634e727d8d9b."

F2. 200 OK sent from AVS to Alice:

5.3. Sending an ARID in a Communication Request

When Alice makes a call to Bob with an ARID, she needs to specify the ARID associated with the URL of the website for validating the ARID in a SIP UA. The SIP UA MUST generate a new SIP header called "Sender-Reference" including a URI, "type", "salt", and "hash_alg" parameters to convey the ARID in the path of an HTTP URL, specify this service, and the salt and the hash algorithm which were used for hashing the querier's contact address described in [Section 5.2](#), respectively. If Alice wants to specify multiple ARIDs, this Sender-References header field includes multiple set of an ARID and related parameters concatenating a comma separator. The SIP UA then sends an SIP INVITE request including the Sender-Reference as shown in the following example. The INVITE request MUST be sent over TLS to protect message confidentiality and integrity.

F3. SIP INVITE from Alice to Bob:

5.4. Validating an ARID to Retrieve User Attributes

When Bob, the recipient of one or more ARIDs, wants to retrieve the caller attributes, the SIP UAS needs to test the validity of the ARIDs on the corresponding AVSes. By prompting Bob or based on his preconfigured information, the SIP UAS first needs to determine whether or not he trusts each domain name of the AVS in the Sender-References header in received SIP INVITE request. Only for trusted AVSes, the SIP UAS looks up the received ARIDs on the corresponding AVSes to retrieve the caller's attributes by using an HTTP GET request as shown in the following example. HTTP messages MUST sent over TLS for security as well as messages between the SIP UAC the AVS. Bob authenticates the AVS using its X.509 PKC delivered in the TLS handshake.

For this validation, the SIP UAS MUST send the ARID found in the Sender-References header field and a hashed querier's contact address generated by the hash algorithm and salt also found in the Sender-References header field. To generate a hashed querier's contact address, the SIP UAS needs to know the original destination address by extracting from the To header or by Bob's pre-configuration especially when he enables call forwarding services. In the following example, the hashed querier's contact address,

"2cf6a1eda3b5205005d25a7d5dcf13bb200fc26a", is generated by `SHA1("dmvb1p03"||"sips:bob@example.com")`. This validation MAY be invoked by a SIP inbound proxy on behalf of the UAS.

F4. HTTP GET from Bob to AVS:

If no AVSes are trusted by Bob, the SIP UAS MUST ignore the Sender-Reference header field and stop any further validation process. If the SIP UAS does not support a hash algorithm specified in the Sender-References header field, or if the SIP UAS does not support the header field, it SHOULD also ignore the header field and continue normal processing of the received SIP request.

If the ARID is valid at the queried time and with the querier's contact address, the AVS MUST respond to the querier with 200 OK in HTTP having the attributes based on the disclosure mode which Alice specifies in the message body, as shown in the following example. The attributes SHOULD be attached as a JSON object or MAY be in XML. If the query is done later than the expiry time, the AVS SHOULD respond with 408 Request Timeout in HTTP. If the querier is not included in the list of desired queriers specified earlier by Alice, the AVS SHOULD respond with 403 Forbidden in HTTP. If the ARID is invalid for other reasons, the AVS MUST respond with 404 Not Found in HTTP.

F5. HTTP 200 OK from AVS to Bob:

If Bob receives a 200 OK in HTTP from the AVS, he is informed that the ARID is valid and attached information is the caller's attributes, for the example above, the caller is a student member in "example.org". With any other responses, Bob knows nothing about the caller's attributes. Based on this information, he determines whether or not to answer the call and adjusts his communication stance accordingly.

6. Sender-References Header Field

The SIP "Sender-References" header field is newly defined to provide the reference information about the sender or the caller. The field consists of one or more sender-ref information. Each sender-ref information consists of three parts: an absolute URI, sender-ref-type, and avs-params. The absolute URI contains the URI of the AVS website including an ARID in the path. The sender-ref-type indicates the service type of using this header field. For this referencing service, it MUST be "avs." The avs-params consists of two parameters: one is to specify a salt and another is for a hash algorithm. Both parameters are used for hashing a contact address to be presented for validation.

The syntax of the Sender-References header field in the ABNF [\[RFC5234\]](#) is as follows:

This Sender-References header field is optionally set in any SIP requests and responses.

[7.](#) Relationship to Existing Mechanisms

This section discusses why this referencing mechanism does not use existing mechanisms that provide an attribute assertion or third-party authentication, such as an X.509 Attribute Certificate (AC), a Security Assertion Markup Language (SAML) assertion, Vouch by Reference [\[RFC5518\]](#), OAuth [\[RFC5849\]](#) or Kerberos [\[RFC4120\]](#). The following table compares our mechanism using an AVS with these existing mechanisms in terms of the trust model they assume, whether or not to need to bind the assertion to the sender ID, and applicable services.

An X.509 Attribute Certificate (AC) provides a superset of features we need for our equivalent trust model. However, an X.509 AC, unlike our mechanism, requires the AC holder information, namely a user's identity, to be bound to the user's attributes. This binding is protected by being digitally signed with the AC issuer's private key. However, the AC issuer does not always have the right to sign the binding since the AC issuer cannot authenticate the user identity issued by a different organization as described in [Section 3.2](#).

Authenticating the user identity requires either the user's PKC or other mechanisms, such as the SIP identity mechanism where the user identity is a user identifier in SIP. These mechanisms are difficult to deploy for each reason. Users' PKCs have not been widely deployed because of difficulty in managing the pair of public and private keys across multiple devices. The sender ID in SIP is usually issued by an administrative domain different from the AC issuer. For these reasons, we need a new mechanism to allow a looser linkage between the sender ID and attributes.

Similar to an X.509 AC, a SAML assertion provides a superset of features we need for our equivalent trust model. Unlike an X.509 AC, a SAML assertion does not require the binding between user attributes and the user identity. Including the user identity into a SAML assertion is optional. To limit queriers, a SAML assertion can restrict its audience by addressing the URIs of specific entities, but they are currently not allowed to address by their hashed names. Thus, with a minor modification in the form of the restricted audience, we can use an XML-based SAML assertion to convey user attributes instead of a JSON object described in [Section 5.4](#). However, using a SAML assertion requires a digital signature by its issuer, which is an application layer protection against message tampering and server impersonation. As discussed in [Section 4.1](#), we prefer a simple transport layer protection to an application layer one, namely protecting a whole message by TLS rather than protecting part of a message by a digital signature. Vouch by Reference (VBR) defines a simple mechanism that vouches a specific type of content claimed by the sender's domain of an email

message. This mechanism uses a new email "VBR-Info" header and a DNS-based server of a third party certification service. If the recipient finds a trusted domain from the certification service providers set in the VBR-info header in a received email message, he looks up an entry of the sender domain on the DNS-based server of a trusted certification service provider. Since VBR assumes the same trust model as ours, it is possible to extend this VBR mechanism to vouch a user's attributes instead of certifying a specific content type for the sender's domain. However, VBR requires the authentication of the sender domain since the server domain is used as a query key. Additionally, it is difficult for a DNS-based server to restrict queriers for each record, mainly private attributes. Consequently, we cannot apply VBR to referencing user attributes.

OAuth is a third-party authentication model for Web services. OAuth uses three tokens to delegate limited permissions of user's resource to another entity called a Consumer. With the OAuth terminology, a caller is a User, the callee is a Consumer, and the AVS is a Service Provider, which is a third-party authenticator. In OAuth, unlike our trust model, the AVS and the callee share a Consumer ID and a key to authenticate the Consumer when the AVS provides one of these three tokens, a Request Token. An unauthorized Request Token is generated upon the Consumer's request. After the Consumer is authorized by a User, it is provided with an authorized Request Token. Since this authorized Request Token has a restricted scope and limited lifetime to access the Users' resources, this Request Token can be used to query the caller's attributes once the caller authorizes that. However, to obtain an authorized Request Token, the callee, who is a Consumer in OAuth, needs to obtain an unauthorized Request Token from the AVS beforehand. To resolve these differences in the trust models and the procedures, it is possible to omit using both an unauthorized and authorized Request Tokens. In addition, using the third token in OAuth, an Access Token, which is exchanged with an authorized Request Token, is not needed. Thus, we do not need unnecessary complexities of using these three types of tokens. Rather, we prefer a simple mechanism, using a single token for SIP.

Kerberos provides strong user-client authentication using a Key Distribution Center (KDC). An Authentication Server in KDC authenticates a user on behalf of a Service Server (SS), and a Ticket Granting Server (TGS) issues a ticket which is effective only for a session between the user and the SS for a limited time period. The Kerberos features cover all our mechanism needs, but the trust model is different. Kerberos assumes that the TGS and the SS share the SS's secret key to allow the SS to verify a received ticket by decrypting with the SS's key without connecting to the TGS. Since using this ticket is a key feature in Kerberos, we cannot omit sharing the SS's key with the TGS to resolve the difference in the trust model. Because of this difference in assumed trust model, we cannot use Kerberos for referencing and validating user attributes.

8. Security Considerations

8.1. Man in the Middle Attacks

Man in the middle attacks need to be prevented on the AVS, connection links, and the recipients of an ARID. To prevent from impersonating a user on the AVS, the AVS MUST authenticate a user using HTTP Basic or Digest authentication, a client X.509 PKC or other mechanisms. To prevent from eavesdropping and message tampering on connection links, all connection links between UAC and the AVS, UAC and UAS, UAS and the AVS MUST be protected using TLS.

To prevent from impersonating user attributes using a stolen ARID, the AVS MUST limit queriers using a specific ARID based on the hashed contact addresses the original requester of the ARID specifies. SIP UAS MUST support SHA1 to hash a contact address. To mitigate the damage of the impersonation in the case where an ARID is stolen with one of these hashed contact addresses, the AVS MUST limit an ARID's lifetime and MAY also limit the number of times it can be resolved.

8.2. Replay Attacks Using a Received ARID

The recipient of an ARID can exploit impersonation just by forwarding a received ARID to another user since this mechanism does not have a tight link between the username of AVS and the caller ID as described in [Section 3.2](#) nor a link between the SIP signaling path and the ARID. To prevent this replay or forwarding attack, the AVS MUST limit queriers for each ARID based on the hashed contact addresses that the original requester of the ARID specifies. This is the same way as preventing impersonation using a stolen ARID described in [Section 8.1](#). Suppose Bob, the recipient of Alice's ARID, forwards the ARID to Carol when sending an instant message to her. In the message, Bob instructs Carol to query his attributes using his hashed contact address, instead of hers. By instructing this wrong way of query, Bob fails in his attempt to masquerade as a user having Alice's attributes. However, despite Alice's original designation, Carol can retrieve Alice's attributes following Bob's wrong instruction, resulting in raising Alice's privacy concerns. This privacy problem is caused by Bob's misbehavior and unavoidable for any attribute mechanisms which others can retrieve the attributes.

8.3. Denial of Service Attacks on the AVS

Another form of potential attacks is denial of service (DoS) attacks by flooding requests to exhaust resources on the AVS. To mitigate the damage from DoS attacks, we need to spare resources for valid requests. For this purpose, the AVS MUST carefully configure TCP and implement user authentication. To detect invalid requests as easily as possible, this mechanism SHOULD use a light query protocol using the RESTful API [\[REST\]](#), which sets a query key in the path of an HTTP URL.

8.4. Phishing Attacks on the AVS

An evil website having a domain name confusingly similar to a well-known AVS makes it possible to steal the password of a user for remote access to the AVS. It is also possible for an evil website to respond to any attribute queries with an HTTP 200 OK response with forged user attributes attached to invalidate the attribute validation service. To prevent these attacks, both a user and the recipient of an ARID MUST use TLS when connecting to the AVS and MUST ensure that the server's PKC has a valid signature for the valid domain name.

9. IANA Considerations

This document defines a new SIP Sender-References header field. This header field needs to be registered by the IANA in the SIP Parameters registry under the Header Fields sub-registry.

10. References

10.1. Normative References

[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.
[RFC2616]	Fielding, R. , Gettys, J. , Mogul, J. , Frystyk, H. , Masinter, L. , Leach, P. and T. Berners-Lee , " Hypertext Transfer Protocol -- HTTP/1.1 ", RFC 2616, June 1999.
[RFC5246]	Dierks, T. and E. Rescorla, " The Transport Layer Security (TLS) Protocol Version 1.2 ", RFC 5246, August 2008.
[RFC5234]	Crocker, D. and P. Overell, " Augmented BNF for Syntax Specifications: ABNF ", STD 68, RFC 5234, January 2008.
[RFC3986]	Berners-Lee, T. , Fielding, R. and L. Masinter , " Uniform Resource Identifier (URI): Generic Syntax ", STD 66, RFC 3986, January 2005.
[SHA1]	National Institute of Science and Technology, "Secure Hash Standard", Federal Information Processing Standard (FIPS) 180-2, August 2002.

10.2. Informative References

, "

[RFC2617]	Franks, J. , Hallam-Baker, P.M. , Hostetler, J.L. , Lawrence, S.D. , Leach, P.J. , Luotonen, A. and L.
-----------	--

	Stewart , " HTTP Authentication: Basic and Digest Access Authentication ", RFC 2617, June 1999.
[RFC4484]	Peterson, J., Polk, J., Sicker, D. and H. Tschofenig, " Trait-Based Authorization Requirements for the Session Initiation Protocol (SIP) ", RFC 4484, August 2006.
[RFC4627]	Crockford, D., " The application/json Media Type for JavaScript Object Notation (JSON) ", RFC 4627, July 2006.
[RFC5280]	Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R. and W. Polk, " Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile ", RFC 5280, May 2008.
[RFC4474]	Peterson, J. and C. Jennings, " Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP) ", RFC 4474, August 2006.
[RFC5755]	Farrell, S., Housley, R. and S. Turner, " An Internet Attribute Certificate Profile for Authorization ", RFC 5755, January 2010.
[RFC5518]	Hoffman, P., Levine, J. and A. Hathcock, " Vouch By Reference ", RFC 5518, April 2009.
[RFC3966]	Schulzrinne, H., " The tel URI for Telephone Numbers ", RFC 3966, December 2004.
[RFC4120]	Neuman, C., Yu, T., Hartman, S. and K. Raeburn, " The Kerberos Network Authentication Service (V5) ", RFC 4120, July 2005.
[RFC5849]	Hammer-Lahav, E., " The OAuth 1.0 Protocol ", RFC 5849, April 2010.
[RFC5751]	Ramsdell, B. and S. Turner, " Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification ", RFC 5751, January 2010.
[AES]	National Institute of Science and Technology, "Specifications for the Advanced Encryption Standard", Federal Information Processing Standard (FIPS) 197, November 2001.
[XML]	Bray, T, Paoli, J, Sperberg-McQueen, C.M, Maler, E and F Yergeau, "Extensible Markup Language (XML) 1.0 (Third Edition)", W3C Recommendation REC-xml-20040204, February 2004.
[SAML]	Security Assertion Markup Language (SAML) V2.0", March 2005.
[REST]	Fielding, R and R Taylor, "Principled design of the modern Web architecture", ACM Transactions on Internet Technology (TOIT) 2-2, May 2002.

[Authors' Addresses](#)

Kumiko Ono Ono Department of Computer Science Columbia University
New York, NY 10027 USA EMail: kumiko@cs.columbia.edu

Henning Schulzrinne Schulzrinne Department of Computer Science
Columbia University New York, NY 10027 USA EMail:
hgs@cs.columbia.edu