

ABFAB
Internet-Draft
Intended status: Standards Track
Expires: July 14, 2016

J. Howlett
Janet
S. Hartman
Painless Security
A. Perez-Mendez, Ed.
University of Murcia
January 11, 2016

A RADIUS Attribute, Binding, Profiles, Name Identifier Format, and
Confirmation Methods for SAML
draft-ietf-abfab-aaa-saml-14

Abstract

This document describes the use of the Security Assertion Mark-up Language (SAML) with RADIUS in the context of the ABFAB architecture. It defines two RADIUS attributes, a SAML binding, a SAML name identifier format, two SAML profiles, and two SAML confirmation methods. The RADIUS attributes permit encapsulation of SAML assertions and protocol messages within RADIUS, allowing SAML entities to communicate using the binding. The two profiles describe the application of this binding for ABFAB authentication and assertion query/request, enabling a Relying Party to request authentication of, or assertions for, users or machines (Clients). These Clients may be named using a NAI name identifier format. Finally, the subject confirmation methods allow requests and queries to be issued for a previously authenticated user or machine without needing to explicitly identify them as the subject. The use of the artifacts defined in this document is not exclusive to ABFAB. They can be applied in any AAA scenario, such as the network access control.

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 July 14, 2016.

Copyright Notice

Copyright (c) 2016 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	3
1.1. Terminology	4
2. Conventions	5
3. RADIUS SAML Attributes	5
3.1. SAML-Assertion attribute	5
3.2. SAML-Protocol attribute	6
4. SAML RADIUS Binding	7
4.1. Required Information	7
4.2. Operation	7
4.3. Processing of names	9
4.3.1. AAA names	9
4.3.2. SAML names	9
4.3.3. Mapping of AAA names in SAML metadata	10
4.3.4. Example of SAML metadata including AAA names	12
4.4. Use of XML Signatures	13
4.5. Metadata Considerations	13
5. Network Access Identifier Name Identifier Format	13
6. RADIUS State Confirmation Method Identifiers	13
7. ABFAB Authentication Profile	14
7.1. Required Information	14
7.2. Profile Overview	14
7.3. Profile Description	16
7.3.1. Client Request to Relying Party	16
7.3.2. Relying Party Issues <samlp:AuthnRequest> to Identity Provider	16
7.3.3. Identity Provider Identifies Client	17
7.3.4. Identity Provider Issues <samlp:Response> to Relying Party	17
7.3.5. Relying Party Grants or Denies Access to Client	17

7.4.	Use of Authentication Request Protocol	17
7.4.1.	<samlp:AuthnRequest> Usage	18
7.4.2.	<samlp:Response> Message Usage	18
7.4.3.	<samlp:Response> Message Processing Rules	19
7.4.4.	Unsolicited Responses	19
7.4.5.	Use of the SAML RADIUS Binding	19
7.4.6.	Use of XML Signatures	20
7.4.7.	Metadata Considerations	20
8.	ABFAB Assertion Query/Request Profile	20
8.1.	Required Information	20
8.2.	Profile Overview	20
8.3.	Profile Description	21
8.3.1.	Differences from the SAML V2.0 Assertion Query/Request Profile	21
8.3.2.	Use of the SAML RADIUS Binding	22
8.3.3.	Use of XML Signatures	22
8.3.4.	Metadata Considerations	22
9.	Privacy considerations	22
10.	Security Considerations	23
11.	IANA Considerations	24
11.1.	RADIUS Attributes	24
11.2.	ABFAB Parameters	24
11.3.	Registration of the ABFAB URN Namespace	25
12.	Acknowledgements	25
13.	References	25
13.1.	Normative References	25
13.2.	Informative References	27
Appendix A.	XML Schema	29
Authors' Addresses	31

1. Introduction

Within the ABFAB (Application Bridging for Federated Access Beyond web) architecture [I-D.ietf-abfab-arch] it is often desirable to convey Security Assertion Mark-up Language (SAML) assertions and protocol messages.

SAML typically only considers the use of HTTP-based transports, known as bindings [OASIS.saml-bindings-2.0-os], which are primarily intended for use with the SAML V2.0 Web Browser Single Sign-On Profile [OASIS.saml-profiles-2.0-os]. However the goal of ABFAB is to extend the applicability of federated identity beyond the Web to other applications by building on the AAA framework. Consequently there exists a requirement for SAML to integrate with the AAA framework and protocols such as RADIUS [RFC2865] and Diameter [RFC6733], in addition to HTTP.

In summary this document specifies:

- o Two RADIUS attributes to encapsulate SAML assertions and protocol messages respectively.
- o A SAML RADIUS binding that defines how SAML assertions and protocol messages can be transported by RADIUS within a SAML exchange.
- o A SAML name identifier format in the form of a Network Access Identifier.
- o A profile of the SAML Authentication Request Protocol that uses the SAML RADIUS binding to effect SAML-based authentication and authorization.
- o A profile of the SAML Assertion Query And Request Protocol that uses the SAML RADIUS binding to effect the query and request of SAML assertions.
- o Two SAML Subject Confirmation Methods for indicating that a user or machine client is the subject of an assertion.

This document adheres to the guidelines stipulated by [OASIS.saml-bindings-2.0-os] and [OASIS.saml-profiles-2.0-os] for defining new SAML bindings and profiles respectively, and other conventions applied formally or otherwise within SAML. In particular, this document provides a 'Required Information' section for the binding and profiles that enumerate:

- o A URI that uniquely identifies the protocol binding or profile.
- o Postal or electronic contact information for the author.
- o A reference to previously defined bindings or profiles that the new binding updates or obsoletes.
- o In the case of a profile, any SAML confirmation method identifiers defined and/or utilized by the profile.

1.1. Terminology

This document uses terminology from a number of related standards, which tend to adopt different terms for similar or identical concepts. In general the document uses, when possible, the ABFAB term for the entity, as described in [I-D.ietf-abfab-arch]. For reference we include this table which maps the different terms into a single view.

Protocol	Client	Relying Party	Identity Provider
ABFAB	Client	Relying Party	Identity Provider
SAML	Subject Principal	Service Provider Requester Consumer	Identity Provider Responder Issuer
RADIUS	User	NAS RADIUS client	AS RADIUS server

Table 1. Terminology

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 RFC 2119 [RFC2119].

3. RADIUS SAML Attributes

The RADIUS SAML binding defined in Section 4 of this document uses two attributes to convey SAML assertions and protocol messages [OASIS.saml-core-2.0-os]. Owing to the typical size of these structures, these attributes use the Long Extended Type format [RFC6929] to encapsulate their data. RADIUS entities MUST NOT include both attributes in the same RADIUS message, as they represent exclusive alternatives to convey SAML information.

3.1. SAML-Assertion attribute

This attribute is used to encode a SAML assertion. The following figure represents the format of this attribute.

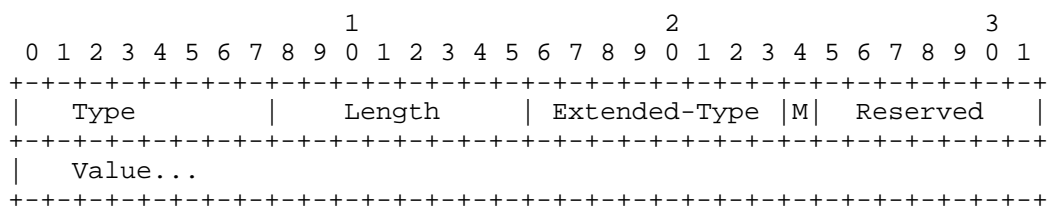


Figure 1: SAML-Assertion format

Type

245 (To be confirmed by IANA)

Length

>= 5

Extended-Type

TBD1

M (More)

As described in [RFC6929].

Reserved

As described in [RFC6929].

Value

One or more octets encoding a SAML assertion.

3.2. SAML-Protocol attribute

This attribute is used to encode a SAML protocol message. The following figure represents the format of this attribute.




Figure 2: SAML-Protocol format

Type

245 (To be confirmed by IANA)

Length

>= 5

Extended-Type

TBD2

M (More)

As described in [RFC6929].

Reserved

As described in [RFC6929].

Value

One or more octets encoding a SAML protocol message.

4. SAML RADIUS Binding

The SAML RADIUS binding defines how RADIUS [RFC2865] can be used to enable a RADIUS client and server to exchange SAML assertions and protocol messages.

4.1. Required Information

Identification: urn:ietf:params:abfab:bindings:radius

Contact information: iesg@ietf.org

Updates: None.

4.2. Operation

In this specification, the Relying Party MUST trust any statement in the SAML messages from the IdP in the same way that it trusts information contained in RADIUS attributes. These entities MUST trust the RADIUS infrastructure to provide integrity of the SAML messages.

Hence, it is REQUIRED that the RADIUS exchange is protected using TLS encryption for RADIUS [RFC6614] to provide confidentiality and integrity protection, unless alternative methods to ensure them are used, such as IPSEC tunnels or a sufficiently secure internal network.

Implementations of this profile can take advantage of mechanisms to permit the transport of longer SAML messages over RADIUS transports, such as the Support of fragmentation of RADIUS packets [RFC7499] or Larger Packets for RADIUS over TCP [I-D.ietf-radext-bigger-packets].

There are two system models for the use of SAML over RADIUS. The first is a request-response model, using the RADIUS SAML-Protocol

attribute defined in Section 3 to encapsulate the SAML protocol messages.

1. The RADIUS client, acting as a Relying Party (RP), transmits a SAML request element within a RADIUS Access-Request message. This message MUST include a single instance of the RADIUS User-Name attribute whose value MUST conform to the Network Access Identifier [RFC7542] scheme. The Relying Party MUST NOT include more than one SAML request element.
2. The RADIUS server, acting as an Identity Provider (IdP), returns a SAML protocol message within a RADIUS Access-Accept or Access-Reject message. These messages necessarily conclude a RADIUS exchange and therefore this is the only opportunity for the Identity Provider to send a response in the context of this exchange. The Identity Provider MUST NOT include more than one SAML response. An IdP that refuses to perform a message exchange with the Relying Party can silently discard the SAML request (this could subsequently be followed by a RADIUS Access-Reject, as the same conditions that cause the IdP to discard the SAML request may also cause the RADIUS server to fail to authenticate).

The second system model permits a RADIUS server acting as an Identity Provider to use the RADIUS SAML-Assertion attribute defined in Section 3 to encapsulate an unsolicited SAML assertion. This attribute MUST be included in a RADIUS Access-Accept message. When included, the attribute MUST contain a single SAML assertion.

RADIUS servers MUST NOT include both the SAML-Protocol and the SAML-Assertion attribute in the same RADIUS message. If an IdP is producing a response to a SAML request, then the first system model is used. An IdP MAY ignore a SAML request and send an unsolicited assertion using the second system model using the RADIUS SAML-Assertion attribute.

In either system model, Identity Providers SHOULD return a RADIUS state attribute as part of the Access-Accept message so that future SAML queries or requests can be run against the same context of an authentication exchange.

This binding is intended to be composed with other uses of RADIUS, such as network access. Therefore, other arbitrary RADIUS attributes MAY be used in either the request or response.

In the case of a SAML processing error, the RADIUS server MAY include a SAML response message with an appropriate value for the `<samlp:Status>` element within the Access-Accept or Access-Reject

packet to notify the client. Alternatively, the RADIUS server can respond without a SAML-Protocol attribute.

4.3. Processing of names

SAML entities using profiles making use of this binding will typically possess both the SAML and AAA names of their correspondents. Frequently these entities will need to apply policies using these names; for example, when deciding to release attributes. Often these policies will be security-sensitive, and so it is important that policy is applied on these names consistently.

4.3.1. AAA names

These rules relate to the processing of AAA names by SAML entities using profiles making use of this binding.

- o Identity Providers SHOULD apply policy based on the Relying Party's identity associated with the RADIUS Access-Request.
- o Relying Parties SHOULD apply policy based on the NAI realm associated with the RADIUS Access-Accept.

4.3.2. SAML names

These rules relate to the processing of SAML names by SAML entities using profiles making use of this binding.

Identity Providers MAY apply policy based on the Relying Party's SAML entityId. In such cases, at least one of the following methods is required in order to establish a relation between the SAML name and the AAA name of the Relying Party:

- o RADIUS client identity in trusted SAML metadata (as described in section Section 4.3.3).
- o RADIUS client identity in trusted digitally signed SAML request.

A digitally signed SAML request without the RADIUS client identity is not sufficient, since a malicious RADIUS entity can observe a SAML message and include it in a different RADIUS message without the consent of the issuer of that SAML message. If an Identity Provider were to process the SAML message without confirming that it applied to the RADIUS message, inappropriate policy would be used.

Relying Parties MAY apply policy based on the SAML issuer's <entityId>. In such cases, at least one of the following methods is

required in order to establish a relationship between the SAML name and the AAA name of the Identity Provider:

- o RADIUS realm in trusted SAML metadata (as described in section Section 4.3.3).
- o RADIUS realm in trusted digitally signed SAML response or assertion.

A digitally signed SAML response alone is not sufficient for the same reasons described above for SAML requests.

4.3.3. Mapping of AAA names in SAML metadata

This section defines extensions to the SAML metadata schema [OASIS.saml-metadata-2.0-os] that are required in order to represent AAA names associated with a particular <EntityDescriptor> element.

In SAML metadata, a single entity may act in many different roles in the support of multiple profiles. This document defines two new roles: RADIUS IDP and RADIUS RP, requiring the declaration of two new subtypes of RoleDescriptorType: RADIUSIDPDescriptorType and RADIUSRPDescriptorType. These subtypes contain the additional elements required to represent AAA names for IDP and RP entities respectively.

4.3.3.1. RADIUSIDPDescriptorType

The RADIUSIDPDescriptorType complex type extends RoleDescriptorType with elements common to IdPs that support RADIUS. It contains the following additional elements:

<RADIUSIDPService> [Zero or More] Zero or more elements of type EndpointType that describe RADIUS endpoints that are associated with the entity.

<RADIUSRealm> [Zero or More] Zero or more elements of type string that represent the acceptable values of the RADIUS realm associated with the entity, obtained from the realm part of RADIUS User-Name attribute.

The following schema fragment defines the RADIUSIDPDescriptorType complex type:

```

    <complexType name="RADIUSIDPDescriptorType">
      <complexContent>
        <extension base="md:RoleDescriptorType">
          <sequence>
            <element ref="abfab:RADIUSIDPService" minOccurs="0" maxOccurs="unbounded"/>
            <element ref="abfab:RADIUSRealm" minOccurs="0" maxOccurs="unbounded"/>
          </sequence>
        </extension>
      </complexContent>
    </complexType>
    <element name="RADIUSIDPService" type="md:EndpointType"/>
    <element name="RADIUSRealm" type="string"/>

```

Figure 3: RADIUSIDPDescriptorType schema

4.3.3.2. RADIUSRPDescriptorType

The RADIUSRPDescriptorType complex type extends RoleDescriptorType with elements common to RPs that support RADIUS. It contains the following additional elements:

- <RADIUSRPService> [Zero or More] Zero or more elements of type EndpointType that describe RADIUS endpoints that are associated with the entity.
- <RADIUSNasIpAddress> [Zero or More] Zero or more elements of type string that represent the acceptable values of the RADIUS NAS-IP-Address or NAS-IPv6-Address attributes associated with the entity.
- <RADIUSNasIdentifier> [Zero or More] Zero or more elements of type string that represent the acceptable values of the RADIUS NAS-Identifier attribute associated with the entity.
- <RADIUSGssEapName> [Zero or More] Zero or more elements of type string that represent the acceptable values of the GSS-EAP acceptor name associated with the entity. The format for this name is described in section 3.1 of [RFC7055], while section 3.4 describes how that name is decomposed and transported using RADIUS attributes.

The following schema fragment defines the RADIUSRPDescriptorType complex type:

```

<complexType name="RADIUSRPDescriptorType">
  <complexContent>
    <extension base="md:RoleDescriptorType">
      <sequence>
        <element ref="md:RADIUSRPService" minOccurs="0" maxOccurs="unbounded" />
        <element ref="md:RADIUSNasIpAddress" minOccurs="0" maxOccurs="unbounded" />
        <element ref="md:RADIUSNasIdentifier" minOccurs="0" maxOccurs="unbounded" />
        <element ref="md:RADIUSGssEapName" minOccurs="0" maxOccurs="unbounded" />
      </sequence>
    </extension>
  </complexContent>
</complexType>
<element name="RADIUSRPService" type="md:EndpointType" />
<element name="RADIUSNasIpAddress" type="string" />
<element name="RADIUSNasIdentifier" type="string" />
<element name="RADIUSGssEapName" type="string" />

```

Figure 4: RADIUSRPDescriptorType schema

4.3.4. Example of SAML metadata including AAA names

The following figures illustrate an example of metadata including AAA names for an IDP and a RP respectively. The IDP's SAML name is "https://IdentityProvider.com/", whereas its RADIUS realm is "idp.com". The RP's SAML name is "https://RelyingParty.com/SAML", being its GSS-EAP acceptor name "nfs/fileserver.rp.com@RP.COM".

```

<EntityDescriptor xmlns="urn:oasis:names:tc:SAML:2.0:metadata"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:abfab="urn:ietf:params:xml:ns:abfab"
  entityID="https://IdentityProvider.com/SAML">
  <RoleDescriptor xsi:type="abfab:RADIUSIDPDescriptorType"
    protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">
    <RADIUSRealm>idp.com</RADIUSRealm>
  </RoleDescriptor>
</EntityDescriptor>

```

Figure 5: Metadata for the IDP

```
<EntityDescriptor xmlns="urn:oasis:names:tc:SAML:2.0:metadata"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:abfab="urn:ietf:params:xml:ns:abfab"
  entityID="https://RelyingParty.com/SAML">
  <RoleDescriptor xsi:type="abfab:RADIUSRPDescriptorType"
    protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">
    <RADIUSGssEapName>nfs/fileserver.rp.com@RP.COM</RADIUSGssEapName>
  </RoleDescriptor>
</EntityDescriptor>
```

Figure 6: Metadata for the RP

4.4. Use of XML Signatures

This binding calls for the use of SAML elements that support XML signatures. To promote interoperability, implementations of this binding **MUST** support a default configuration that does not require the use of XML signatures. Implementations **MAY** choose to use XML signatures.

4.5. Metadata Considerations

These binding and profiles are mostly intended to be used without metadata. In this usage, RADIUS infrastructure is used to provide integrity and naming of the SAML messages and assertions. RADIUS configuration is used to provide policy, including which attributes are accepted from a Relying Party and which attributes are sent by an Identity Provider.

Nevertheless, if metadata is used, the roles describe in section Section 4.3.3 **MUST** be present.

5. Network Access Identifier Name Identifier Format

URI: urn:ietf:params:abfab:nameid-format:nai

Indicates that the content of the element is in the form of a Network Access Identifier (NAI) using the syntax described by [RFC7542].

6. RADIUS State Confirmation Method Identifiers

URI: urn:ietf:params:abfab:cm:user

URI: urn:ietf:params:abfab:cm:machine

Indicates that the Subject is the system entity (either the user or machine) authenticated by a previously transmitted RADIUS Access-

Accept message, as identified by the value of that RADIUS message's State attribute.

7. ABFAB Authentication Profile

In the scenario supported by the ABFAB Authentication Profile, a Client controlling a User Agent requests access to a Relying Party. The Relying Party uses RADIUS to authenticate the Client. In particular, the Relying Party, acting as a RADIUS client, attempts to validate the Client's credentials against a RADIUS server acting as the Client's Identity Provider. If the Identity Provider successfully authenticates the Client, it produces an authentication assertion which is consumed by the Relying Party. This assertion MAY include a name identifier that can be used between the Relying Party and the Identity Provider to refer to the Client.

7.1. Required Information

Identification: urn:ietf:params:abfab:profiles:authentication

Contact information: iesg@ietf.org

SAML Confirmation Method Identifiers: The SAML V2.0 "RADIUS State" confirmation method identifiers, either urn:ietf:params:abfab:cm:user or urn:ietf:params:abfab:cm:machine, are used by this profile.

Updates: None.

7.2. Profile Overview

To implement this scenario, this profile of the SAML Authentication Request protocol MUST be used in conjunction with the SAML RADIUS binding defined in Section 4.

This profile is based on the SAML V2.0 Web Browser Single Sign-On Profile [OASIS.saml-profiles-2.0-os]. There are some important differences, specifically:


Authentication: This profile does not require the use of any particular authentication method. The ABFAB architecture does require the use of EAP [RFC3579], but this specification may be used in other non-ABFAB scenarios.

Bindings: This profile does not use HTTP-based bindings. Instead all SAML protocol messages are transported using the SAML RADIUS binding defined in Section 4. This is intended to reduce the number of bindings that implementations must support to be interoperable.

Requests: The profile does not permit the Relying Party to name the `<saml:Subject>` of the `<samlp:AuthnRequest>`. This is intended to simplify implementation and interoperability.

Responses: The profile only permits the Identity Provider to return a single SAML message or assertion that **MUST** contain exactly one authentication statement. Other statements may be included within this assertion at the discretion of the Identity Provider. This is intended to simplify implementation and interoperability.

Figure 7 below illustrates the flow of messages within this profile.



The following steps are described by the profile. Within an individual step, there may be one or more actual message exchanges.

Figure 7

1. Client request to Relying Party (Section 7.3.1): In step 1, the Client, via a User Agent, makes a request for a secured resource at the Relying Party. The Relying Party determines that no security context for the Client exists and initiates the authentication process.
2. Relying Party issues `<samlp:AuthnRequest>` to Identity Provider (Section 7.3.2). In step 2, the Relying Party may optionally issue a `<samlp:AuthnRequest>` message to be delivered to the Identity Provider using the SAML-Protocol RADIUS attribute.

3. Identity Provider identifies Client (Section 7.3.3). In step 3, the Client is authenticated and identified by the Identity Provider, while honoring any requirements imposed by the Relying Party in the <samlp:AuthnRequest> message if provided.
4. Identity Provider issues <samlp:Response> to Relying Party (Section 7.3.4). In step 4, the Identity Provider issues a <samlp:Response> message to the Relying Party using the SAML RADIUS binding. The response either indicates an error or includes a SAML Authentication Statement in exactly one SAML Assertion. If the RP did not send an <samlp:AuthnRequest>, the IdP issues an unsolicited <samlp:Assertion>, as described in Section 7.4.4.
5. Relying Party grants or denies access to Client (Section 7.3.5). In step 5, having received the response from the Identity Provider, the Relying Party can respond to the Client with its own error, or can establish its own security context for the Client and return the requested resource.

7.3. Profile Description

The ABFAB Authentication Profile is a profile of the SAML V2.0 Authentication Request Protocol [OASIS.saml-core-2.0-os]. Where both specifications conflict, the ABFAB Authentication Profile takes precedence.

7.3.1. Client Request to Relying Party

The profile is initiated by an arbitrary Client request to the Relying Party. There are no restrictions on the form of the request. The Relying Party is free to use any means it wishes to associate the subsequent interactions with the original request. The Relying Party, acting as a RADIUS client, attempts to authenticate the Client.

7.3.2. Relying Party Issues <samlp:AuthnRequest> to Identity Provider

The Relying Party uses RADIUS to communicate with the Client's Identity Provider. The Relying Party MAY include a <samlp:AuthnRequest> within this RADIUS Access-Request message using the SAML-Protocol RADIUS attribute. The next hop destination MAY be the Identity Provider or alternatively an intermediate RADIUS proxy.

Profile-specific rules for the contents of the <samlp:AuthnRequest> element are given in Section 7.4.1.

7.3.3. Identity Provider Identifies Client

The Identity Provider MUST establish the identity of the Client using a RADIUS authentication method, or else it will return an error. If the ForceAuthn attribute on the <samlp:AuthnRequest> element (if sent by the Relying Party) is present and true, the Identity Provider MUST freshly establish this identity rather than relying on any existing session state it may have with the Client (for example, TLS state that may be used for session resumption). Otherwise, and in all other respects, the Identity Provider may use any method to authenticate the Client, subject to the constraints called out in the <samlp:AuthnRequest> message.

7.3.4. Identity Provider Issues <samlp:Response> to Relying Party

The Identity Provider MUST conclude the authentication in a manner consistent with the RADIUS authentication result. The IdP MAY issue a <samlp:Response> message to the Relying Party that is consistent with the authentication result, as described in [OASIS.saml-core-2.0-os]. This SAML response is delivered to the Relying Party using the SAML RADIUS binding described in Section 4.

Profile-specific rules regarding the contents of the <samlp:Response> element are given in Section 7.4.2.

7.3.5. Relying Party Grants or Denies Access to Client

If a <samlp:Response> message is issued by the Identity Provider, the Relying Party MUST process that message and any enclosed assertion elements as described in [OASIS.saml-core-2.0-os]. Any subsequent use of the assertion elements is at the discretion of the Relying Party, subject to any restrictions contained within the assertions themselves or from any previously established out-of-band policy that governs the interaction between the Identity Provider and the Relying Party.

7.4. Use of Authentication Request Protocol

This profile is based on the Authentication Request Protocol defined in [OASIS.saml-core-2.0-os]. In the nomenclature of actors enumerated in section 3.4 of that document, the Relying Party is the requester, the User Agent is the attesting entity and the Client is the Requested Subject.

7.4.1. <samlp:AuthnRequest> Usage

The Relying Party MUST NOT include a <saml:Subject> element in the request. The authenticated RADIUS identity identifies the Client to the Identity Provider.

A Relying Party MAY include any message content described in [OASIS.saml-core-2.0-os], section 3.4.1. All processing rules are as defined in [OASIS.saml-core-2.0-os].

If the Relying Party wishes to permit the Identity Provider to establish a new identifier for the Client if none exists, it MUST include a <saml:NameIDPolicy> element with the AllowCreate attribute set to "true". Otherwise, only a Client for whom the Identity Provider has previously established an identifier usable by the Relying Party can be authenticated successfully.

The <samlp:AuthnRequest> message MAY be signed. Authentication and integrity are also provided by the SAML RADIUS binding.

7.4.2. <samlp:Response> Message Usage

If the Identity Provider cannot or will not satisfy the request, it MUST either respond with a <samlp:Response> message containing an appropriate error status code or codes and/or respond with a RADIUS Access-Reject message.

If the Identity Provider wishes to return an error, it MUST NOT include any assertions in the <samlp:Response> message. Otherwise, if the request is successful (or if the response is not associated with a request), the <samlp:Response> element is subject to the following constraints:

- o It MAY be signed.
- o It MUST contain exactly one assertion. The <saml:Subject> element of this assertion MUST refer to the authenticated RADIUS user.
- o The assertion MUST contain a <saml:AuthnStatement>. Besides, the assertion MUST contain a <saml:Subject> element with at least one <saml:SubjectConfirmation> element containing a Method of urn:ietf:params:abfab:cm:user or urn:ietf:params:abfab:cm:machine that reflects the authentication of the Client to the Identity Provider. Since the containing message is in response to an <samlp:AuthnRequest>, the InResponseTo attribute (both in the <saml:SubjectConfirmationData> and in the <saml:Response> elements) MUST match the request's ID. The <saml:Subject> element

MAY use the NAI Name Identifier Format described in Section 5 to establish an identifier between the Relying Party and the IdP.

- o Other conditions MAY be included as requested by the Relying Party or at the discretion of the Identity Provider. The Identity Provider is NOT obligated to honor the requested set of conditions in the <samlp:AuthnRequest>, if any.

7.4.3. <samlp:Response> Message Processing Rules

The Relying Party MUST do the following:

- o Assume that the Client's identifier implied by a SAML <Subject> element, if present, takes precedence over an identifier implied by the RADIUS User-Name attribute.
- o Verify that the InResponseTo attribute in the "RADIUS State" <saml:SubjectConfirmationData> equals the ID of its original <samlp:AuthnRequest> message, unless the response is unsolicited, in which case the attribute MUST NOT be present.
- o If a <saml:AuthnStatement> used to establish a security context for the Client contains a SessionNotOnOrAfter attribute, the security context SHOULD be discarded once this time is reached, unless the Relying Party reestablishes the Client's identity by repeating the use of this profile.
- o Verify that any assertions relied upon are valid according to processing rules in [OASIS.saml-core-2.0-os].
- o Any assertion which is not valid, or whose subject confirmation requirements cannot be met MUST be discarded and MUST NOT be used to establish a security context for the Client.

7.4.4. Unsolicited Responses

An Identity Provider MAY initiate this profile by delivering an unsolicited assertion to a Relying Party. This MUST NOT contain any <saml:SubjectConfirmationData> elements containing an InResponseTo attribute.

7.4.5. Use of the SAML RADIUS Binding

It is RECOMMENDED that the RADIUS exchange is protected using TLS encryption for RADIUS [RFC6614] to provide confidentiality and integrity protection.

7.4.6. Use of XML Signatures

This profile calls for the use of SAML elements that support XML signatures. To promote interoperability implementations of this profile **MUST NOT** require the use of XML signatures. Implementations **MAY** choose to use XML signatures.

7.4.7. Metadata Considerations

There are no metadata considerations particular to this profile, aside from those applying to the use of the RADIUS binding.

8. ABFAB Assertion Query/Request Profile

This profile builds on the SAML V2.0 Assertion Query/Request Profile defined by [OASIS.saml-profiles-2.0-os]. That profile describes the use of the Assertion Query and Request Protocol defined by section 3.3 of [OASIS.saml-core-2.0-os] with synchronous bindings, such as the SOAP binding defined in [OASIS.saml-bindings-2.0-os].

While the SAML V2.0 Assertion Query/Request Profile is independent of the underlying binding, it is nonetheless useful to describe the use of the SAML RADIUS binding defined in Section 4 of this document, in the interests of promoting interoperable implementations, particularly as the SAML V2.0 Assertion Query/Request Profile is most frequently discussed and implemented in the context of the SOAP binding.

8.1. Required Information

Identification: urn:ietf:params:abfab:profiles:query

Contact information: iesg@ietf.org

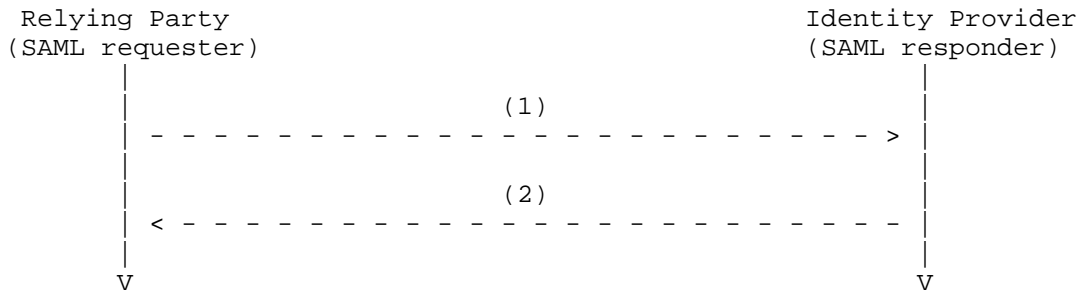
Description: Given below.

Updates: None.

8.2. Profile Overview

As with the SAML V2.0 Assertion Query/Request Profile defined by [OASIS.saml-profiles-2.0-os] the message exchange and basic processing rules that govern this profile are largely defined by Section 3.3 of [OASIS.saml-core-2.0-os] that defines the messages to be exchanged, in combination with the binding used to exchange the messages. The SAML RADIUS binding described in this document defines the binding of the message exchange to RADIUS. Unless specifically noted here, all requirements defined in those specifications apply.

Figure 8 below illustrates the basic template for the query/request profile.



The following steps are described by the profile.

Figure 8

1. Query/Request issued by Relying Party: In step 1, a Relying Party initiates the profile by sending an <AssertionIDRequest>, <SubjectQuery>, <AuthnQuery>, <AttributeQuery>, or <AuthzDecisionQuery> message to a SAML authority.
2. <Response> issued by SAML Authority: In step 2, the responding SAML authority (after processing the query or request) issues a <Response> message to the Relying Party.

8.3. Profile Description

8.3.1. Differences from the SAML V2.0 Assertion Query/Request Profile

This profile is identical to the SAML V2.0 Assertion Query/Request Profile, with the following exceptions:

- o When processing the SAML request, the IdP MUST give precedence to the Client's identifier implied by RADIUS State attribute, if present, over the identifier implied by the SAML request's <Subject>, if any.
- o In respect to sections 6.3.1 and 6.5 of [OASIS.saml-profiles-2.0-os], this profile does not consider the use of metadata (as in [OASIS.saml-metadata-2.0-os]). See Section 8.3.4.
- o In respect to sections 6.3.2, 6.4.1, and 6.4.2 of [OASIS.saml-profiles-2.0-os], this profile additionally stipulates that implementations of this profile MUST NOT require the use of XML signatures. See Section 8.3.3.

8.3.2. Use of the SAML RADIUS Binding

The RADIUS Access-Request sent by the Relying Party:

- o MUST include an instance of the RADIUS Service-Type attribute, having a value of Authorize-Only.
- o SHOULD include the RADIUS State attribute, where this Query/Request pertains to previously authenticated Client.

When processing the SAML request, the IdP MUST give precedence to the Client's identifier implied by RADIUS State attribute over the identifier implied by the SAML request's <Subject>, if any.

It is RECOMMENDED that the RADIUS exchange is protected using TLS encryption for RADIUS [RFC6614] to provide confidentiality and integrity protection.

8.3.3. Use of XML Signatures

This profile calls for the use of SAML elements that support XML signatures. To promote interoperability implementations of this profile MUST NOT require the use of XML signatures. Implementations MAY choose to use XML signatures.

8.3.4. Metadata Considerations

There are no metadata considerations particular to this profile, aside from those applying to the use of the RADIUS binding.

9. Privacy considerations

The profiles defined in this document allow a Relying Party to request specific information about the Client, and allow an IdP to disclose information about that Client. In this sense, Identity Providers MUST apply policy to decide what information is released to a particular Relying Party. Moreover, the identity of the Client is typically hidden from the Relying Party unless informed by the Identity Provider. Conversely, the Relying Party does typically know the realm of the IdP, as it is required to route the RADIUS packets to the right destination.

The kind of information that is released by the IdP can include generic attributes such as affiliation shared by many Clients. But even these generic attributes can help to identify a specific Client. Other kinds of attributes may also provide a Relying Party with the ability to link the same Client between different sessions. Finally, other kind of attributes might provide a group of Relying Parties

with the ability to link the Client between them or with personally identifiable information about the Client.

These profiles do not directly provide a Client with a mechanism to express preferences about what information is released. That information can be expressed out-of-band, for example as part of the enrollment process.

The Relying Party may disclose privacy-sensitive information about itself as part of the request, although this is unlikely in typical deployments.

If RADIUS proxies are used and encryption is not used, the attributes disclosed by the IdP are visible to the proxies. This is a significant privacy exposure in some deployments. Ongoing work is exploring mechanisms for creating TLS connections directly between the RADIUS client and the RADIUS server to reduce this exposure. If proxies are used, the impact of exposing SAML assertions to the proxies needs to be carefully considered.

The use of TLS to provide confidentiality for the RADIUS exchange is strongly encouraged. Without this, passive eavesdroppers can observe the assertions.

10. Security Considerations

In this specification, the Relying Party **MUST** trust any statement in the SAML messages from the IdP in the same way that it trusts information contained in RADIUS attributes. These entities **MUST** trust the RADIUS infrastructure to provide integrity of the SAML messages.

Furthermore, the Relying Party **MUST** apply policy and filter the information based on what information the IdP is permitted to assert and on what trust is reasonable to place in proxies between them.

XML signatures and encryption are provided as an **OPTIONAL** mechanism for end-to-end security. These mechanism can protect SAML messages from being modified by proxies in the RADIUS infrastructure. These mechanisms are not mandatory-to-implement. It is believed that ongoing work to provide direct TLS connections between a RADIUS client and RADIUS server will provide similar assurances but better deployability. XML security is appropriate for deployments where end-to-end security is required but proxies cannot be removed or where SAML messages need to be verified at a later time or by parties not involved in the authentication exchange.

11. IANA Considerations

11.1. RADIUS Attributes

The authors request that Attribute Types and Attribute Values defined in this document be registered by the Internet Assigned Numbers Authority (IANA) from the RADIUS namespaces as described in the "IANA Considerations" section of [RFC3575], in accordance with BCP 26 [RFC5226]. For RADIUS packets, attributes and registries created by this document IANA is requested to place them at <http://www.iana.org/assignments/radius-types>.

In particular, this document defines two new RADIUS attributes, entitled "SAML-Assertion" and "SAML-Protocol" (see Section 3), with assigned values of 245.TBD1 and 245.TBD2 from the Long Extended Space of [RFC6929]:

Type	Ext. Type	Name	Length	Meaning
----	-----	-----	-----	-----
245	TBD1	SAML-Assertion	>=5	Encodes a SAML assertion
245	TBD2	SAML-Protocol	>=5	Encodes a SAML protocol message

11.2. ABFAB Parameters

A new top-level registry is created titled "ABFAB Parameters".

In this top-level registry, a sub-registry titled "ABFAB URN Parameters" is created. Registration in this registry is by the IETF review or expert review procedures [RFC5226].

This paragraph gives guidance to designated experts. Registrations in this registry are generally only expected as part of protocols published as RFCs on the IETF stream; other URIs are expected to be better choices for non-IETF work. Expert review is permitted mainly to allow early registration related to specifications under development when the community believes they have reached sufficient maturity. The expert SHOULD evaluate the maturity and stability of such an IETF-stream specification. Experts SHOULD review anything not from the IETF stream for consistency and consensus with current practice. Today such requests would not typically be approved.

If a parameter named "paramname" is to be registered in this registry, then its URN will be "urn:ietf:params:abfab:paramname". The initial registrations are as follows:

Parameter	Reference
bindings:radius	Section 4
nameid-format:nai	Section 5
profiles:authentication	Section 7
profiles:query	Section 8
cm:user	Section 6
cm:machine	Section 6

ABFAB Parameters

11.3. Registration of the ABFAB URN Namespace

IANA is requested to register the "abfab" URN sub-namespace in the IETF URN sub-namespace for protocol parameters defined in [RFC3553].

Registry Name: abfab

Specification: draft-ietf-abfab-aaa-saml

Repository: ABFAB URN Parameters (Section Section 11.2)

Index Value: Sub-parameters MUST be specified in UTF-8 using standard URI encoding where necessary.

12. Acknowledgements

The authors would like to acknowledge the OASIS Security Services (SAML) Technical Committee, and Scott Cantor in particular, for their help with the SAML-related material.

The authors would also like to acknowledge the collaboration of Jim Schaad, Leif Johansson, Klaas Wierenga, Stephen Farrell, Gabriel Lopez, and Rafael Marin, who have provided valuable comments on this document.

13. References

13.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

- [RFC2865] Rigney, C., Willens, S., Rubens, A., and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, DOI 10.17487/RFC2865, June 2000, <<http://www.rfc-editor.org/info/rfc2865>>.
- [RFC3579] Aboba, B. and P. Calhoun, "RADIUS (Remote Authentication Dial In User Service) Support For Extensible Authentication Protocol (EAP)", RFC 3579, DOI 10.17487/RFC3579, September 2003, <<http://www.rfc-editor.org/info/rfc3579>>.
- [RFC6614] Winter, S., McCauley, M., Venaas, S., and K. Wierenga, "Transport Layer Security (TLS) Encryption for RADIUS", RFC 6614, DOI 10.17487/RFC6614, May 2012, <<http://www.rfc-editor.org/info/rfc6614>>.
- [RFC6929] DeKok, A. and A. Lior, "Remote Authentication Dial In User Service (RADIUS) Protocol Extensions", RFC 6929, DOI 10.17487/RFC6929, April 2013, <<http://www.rfc-editor.org/info/rfc6929>>.
- [RFC3575] Aboba, B., "IANA Considerations for RADIUS (Remote Authentication Dial In User Service)", RFC 3575, DOI 10.17487/RFC3575, July 2003, <<http://www.rfc-editor.org/info/rfc3575>>.
- [RFC7542] DeKok, A., "The Network Access Identifier", RFC 7542, DOI 10.17487/RFC7542, May 2015, <<http://www.rfc-editor.org/info/rfc7542>>.
- [OASIS.saml-bindings-2.0-os]
Cantor, S., Hirsch, F., Kemp, J., Philpott, R., and E. Maler, "Bindings for the OASIS Security Assertion Markup Language (SAML) V2.0", OASIS Standard saml-bindings-2.0-os, March 2005.
- [OASIS.saml-core-2.0-os]
Cantor, S., Kemp, J., Philpott, R., and E. Maler, "Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V2.0", OASIS Standard saml-core-2.0-os, March 2005.
- [OASIS.saml-profiles-2.0-os]
Hughes, J., Cantor, S., Hodges, J., Hirsch, F., Mishra, P., Philpott, R., and E. Maler, "Profiles for the OASIS Security Assertion Markup Language (SAML) V2.0", OASIS Standard OASIS.saml-profiles-2.0-os, March 2005.

[OASIS.saml-metadata-2.0-os]

Cantor, S., Moreh, J., Philpott, R., and E. Maler,
"Metadata for the Security Assertion Markup Language
(SAML) V2.0", OASIS Standard saml-metadata-2.0-os, March
2005.

13.2. Informative References

- [RFC3553] Mealling, M., Masinter, L., Hardie, T., and G. Klyne, "An IETF URN Sub-namespace for Registered Protocol Parameters", BCP 73, RFC 3553, DOI 10.17487/RFC3553, June 2003, <<http://www.rfc-editor.org/info/rfc3553>>.
- [RFC6733] Fajardo, V., Ed., Arkko, J., Loughney, J., and G. Zorn, Ed., "Diameter Base Protocol", RFC 6733, DOI 10.17487/RFC6733, October 2012, <<http://www.rfc-editor.org/info/rfc6733>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC7055] Hartman, S., Ed. and J. Howlett, "A GSS-API Mechanism for the Extensible Authentication Protocol", RFC 7055, DOI 10.17487/RFC7055, December 2013, <<http://www.rfc-editor.org/info/rfc7055>>.
- [RFC7499] Perez-Mendez, A., Ed., Marin-Lopez, R., Pereniguez-Garcia, F., Lopez-Millan, G., Lopez, D., and A. DeKok, "Support of Fragmentation of RADIUS Packets", RFC 7499, DOI 10.17487/RFC7499, April 2015, <<http://www.rfc-editor.org/info/rfc7499>>.
- [I-D.ietf-abfab-arch]
Howlett, J., Hartman, S., Tschofenig, H., Lear, E., and J. Schaad, "Application Bridging for Federated Access Beyond Web (ABFAB) Architecture", draft-ietf-abfab-arch-13 (work in progress), July 2014.
- [I-D.ietf-radext-bigger-packets]
Hartman, S., "Larger Packets for RADIUS over TCP", draft-ietf-radext-bigger-packets-05 (work in progress), December 2015.

[W3C.REC-xmlschema-1]

Thompson, H., Beech, D., Maloney, M., and N. Mendelsohn,
"XML Schema Part 1: Structures", W3C REC-xmlschema-1, May
2001, <<http://www.w3.org/TR/xmlschema-1/>>.

Appendix A. XML Schema

The following schema formally defines the "urn:ietf:params:xml:ns:abfab" namespace used in this document, in conformance with [W3C.REC-xmlschema-1] While XML validation is optional, the schema that follows is the normative definition of the constructs it defines. Where the schema differs from any prose in this specification, the schema takes precedence.

```

<schema
  targetNamespace="urn:ietf:params:xml:ns:abfab"
  xmlns="http://www.w3.org/2001/XMLSchema"
  xmlns:md="urn:oasis:names:tc:SAML:2.0:metadata"
  xmlns:abfab="urn:ietf:params:xml:ns:abfab"
  elementFormDefault="unqualified"
  attributeFormDefault="unqualified"
  blockDefault="substitution"
  version="1.0">

  <import namespace="urn:oasis:names:tc:SAML:2.0:metadata"/>

  <complexType name="RADIUSIDPDescriptorType">
    <complexContent>
      <extension base="md:RoleDescriptorType">
        <sequence>
          <element ref="abfab:RADIUSIDPService" minOccurs="0" maxOccurs="unbounded"/>
          <element ref="abfab:RADIUSRealm" minOccurs="0" maxOccurs="unbounded"/>
        </sequence>
      </extension>
    </complexContent>
  </complexType>
  <element name="RADIUSIDPService" type="md:EndpointType"/>
  <element name="RADIUSRealm" type="string"/>

  <complexType name="RADIUSRPDescriptorType">
    <complexContent>
      <extension base="md:RoleDescriptorType">
        <sequence>
          <element ref="md:RADIUSRPService" minOccurs="0" maxOccurs="unbounded"/>
          <element ref="md:RADIUSNasIpAddress" minOccurs="0" maxOccurs="unbounded"/>
          <element ref="md:RADIUSNasIdentifier" minOccurs="0" maxOccurs="unbounded"/>
          <element ref="md:RADIUSGssEapName" minOccurs="0" maxOccurs="unbounded"/>
        </sequence>
      </extension>
    </complexContent>
  </complexType>
  <element name="RADIUSRPService" type="md:EndpointType"/>
  <element name="RADIUSNasIpAddress" type="string"/>
  <element name="RADIUSNasIdentifier" type="string"/>
  <element name="RADIUSGssEapName" type="string"/>

</schema>

```

Authors' Addresses

Josh Howlett
Janet
Lumen House, Library Avenue, Harwell
Oxford OX11 0SG
UK

Phone: +44 1235 822363
EMail: Josh.Howlett@ja.net

Sam Hartman
Painless Security

EMail: hartmans-ietf@mit.edu

Alejandro Perez-Mendez (editor)
University of Murcia
Campus de Espinardo S/N, Faculty of Computer Science
Murcia 30100
Spain

Phone: +34 868 88 46 44
EMail: alex@um.es

ABFAB
Internet-Draft
Intended status: Informational
Expires: September 22, 2016

R. Smith
Cardiff University
M. Donnelly
Painless Security
March 21, 2016

Application Bridging for Federated Access Beyond web (ABFAB) Usability
and User Interface Considerations
draft-ietf-abfab-usability-ui-considerations-04

Abstract

The real world use of ABFAB-based technologies requires that any identity that is to be used for authentication has to be configured on the ABFAB-enabled client device. Achieving this requires software on that device (either built into the operating system or a standalone utility) that will interact with the user, managing their identity information and identity-to-service mappings. All designers of software to fulfil this role will face the same set of challenges. This document aims to document these challenges with the aim of producing well-thought out UIs with some degree of consistency between implementations.

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 September 22, 2016.

Copyright Notice

Copyright (c) 2016 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	3
2. Conventions	3
3. Terminology	3
4. Context	4
5. Considerations around Terminology	5
5.1. Identity	6
5.2. Services	6
5.3. Identity to Service Mapping	6
6. Considerations around Management of Identities	7
6.1. Information associated with each Identity	7
6.2. Information associated with each Identity Provider	8
6.3. Storage of Identity Information	9
6.4. Adding/Association of an Identity	10
6.4.1. Identity Provider Addition	10
6.4.2. Identity Addition	12
6.5. Modifying Identity Information	14
6.5.1. Manual Modification	14
6.5.2. Automated Modification	15
6.6. Verifying an identity	15
6.7. Removing an Identity	15
6.7.1. Manual Removal	15
6.7.2. Automated Removal	16
6.8. Storing an Identity with or without credentials	16
7. Considerations around Management of Service to Identity Mappings	16
7.1. Associating a Service with an Identity	17
7.1.1. User-driven Manual Association	17
7.1.2. Automated Rules-based Association	17
7.1.3. Association Conflicts	17
7.2. Disassociating a Service with an Identity	18
7.3. Listing Services and Identities	19
7.4. Showing the Service that is requesting Authentication	19
7.5. Showing the Identity currently in use	19
7.6. Multiple Identities for a Particular Service	20
7.7. Not using ABFAB for a Particular Service	20
8. Handling of Errors	20
8.1. Errors in GSS-API	20
8.1.1. Log of Errors	21

8.2. Examples of errors	21
9. Handling of Successes	22
9.1. Reporting Authentication Success on First Use of Identity	22
9.2. Reporting Authentication Success	22
10. Other Considerations	22
10.1. Identity Selector Taking Focus	22
10.2. Import/Export of Credentials	22
11. Security Considerations	23
12. Privacy Considerations	24
13. IANA Considerations	24
14. Contributors	25
15. Acknowledgements	25
16. References	25
16.1. Normative References	25
16.2. Informative References	26
Appendix A. Change Log	27
Appendix B. Open Issues	29
Authors' Addresses	29

1. Introduction

The use of ABFAB-based technologies requires that any identity that is to be used for authentication has to be configured on the client device. Achieving this requires software on that device (either built into the operating system or a standalone utility) that will interact with the user, and manage the user's identities and credential-to-service mappings. Anyone designing that software will face the same set of challenges.

This document does not intend to supplant evidence-based UI design guidelines; implementers of identity selectors are strongly encouraged to understand the latest in HCI and UX thought and practice. Instead, it aims to document the common challenges faced by implementers with the aim of providing a common starting point for implementers in the hope that this aids in producing well-thought out UIs with some degree of consistency.

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 [RFC2119].

3. Terminology

Various items of terminology used in the document are heavily overloaded in that they mean a variety of different things to different people. In an attempt to minimise this problem, this

section gives a brief description of the main items of terminology used in order to aid a consistent understanding of this document.

- o NAI: Network Access Identifier - a standard way of identifying a user and assisting in the routing of an authentication request (see [RFC4282]).
- o Identity: In this context, an identity is a credential given to a user by a particular organisation with which they have an association. A user may have multiple identities - potentially multiple identities per organisation, and also across multiple organisations. Each identity will consist of an NAI, alongside other information that supports authentication. Note that in other contexts the usual use of "identity" would match our use of "user", whereas the usual use of "identifier" matches our use of identity.
- o Service: The thing that the user is attempting to authenticate to via ABFAB technology. See [I-D.ietf-abfab-usecases] for some example ABFAB use cases. Also known as the Relying Party.
- o Identity Provider: The thing able to make access management decisions about the Identity.
- o Identity Selector: A piece of software that enables the process by which the GSS-API acquires the identity to use with a particular service. An Identity Selector typically would allow the user to configure a set of identities along with service to identity mappings.
- o Trust anchor: An authoritative source of verification of a particular ABFAB Identity Provider, used to allow authentication of an Identity Provider using X.509 [RFC5280]. Typically this will be a commercial CA to allow authentication via chain of trust, or a preconfigured non-commercial certificate (e.g. self-signed).
- o Credential: Whatever is used by the user to authenticate themselves with a particular NAI. What exactly this will be will be dependent on the EAP method being used, but is likely to be something like a password or a certificate.

4. Context

When using the ABFAB architecture (see [I-D.ietf-abfab-arch]) to perform federated authentication to some service, a user will need to provide identity information that they wish to use to authenticate to that particular service. This will happen through a process of the

application calling the GSS-API, which will in turn gather the user's credentials through some process. We will call this process the "identity selector" in this document (though note that this is not a recommendation on terminology for the process).

The simplest way to achieve the desired effect would be a process that simply takes the credentials from the currently logged in user (e.g. the Windows Domain Credentials) and uses those for all services that request authenticate through ABFAB. This approach gives ultimate simplicity in terms of UI (it wouldn't have one) but the least flexibility (the user has to use a single identity for everything). If there is ever to be a requirement for a user to use a different set of credentials for a service, or a requirement for the user to use ABFAB to authenticate to the operating system, then something more complex will be needed.

Where there is a requirement for multiple credentials to be supported, there are at least two methods that could be employed to configure identities and associated information:

- o They could be configured manually by the user in a configuration file that could be edited by hand or some such simple process, and read by the GSS-API mechanism. While this could work very well functionally, in practice only a small subset of users would be happy with - and able to - configure their identities in such a manner.
- o They could be configured through some interactive process. For ease of use this should have a simple UI, although to support some use cases a headless mode (i.e. a way of interacting with the identity selector when there is no GUI present) may need to be supported.

When designing an identity selector with a UI (or indeed, with a headless mode), any implementer will share a common set of usability considerations inherent to the context. This document aims to explore these considerations, and provide advice and guidance on addressing them where possible.

5. Considerations around Terminology

Anyone designing an identity selector will have to grapple with choosing terminology that the average user has some chance of understanding. This terminology can split into a few main functional areas, as discussed next.

5.1. Identity

The first area where terminology is needed is around the identity/identities of the user. Users are typically used to seeing a variety of terms for aspects of their identity in the federated sense, and an even larger variety in the wider Internet sense. For example, in the federated sense some of these terms include "username", "login", "network account", "institutional account", "home organisation account", "credentials", and a myriad of other such terms. However, NAI - the technically correct name for their identity in an ABFAB sense - is highly unlikely to be one of these terms that users are used to seeing. Further, given that the NAI superficially looks like an email address, there is a definite potential for confusion.

Implementers of an identity selector will need to carefully consider their intended audience for both their level of technical capability and the existing terminology that they may have been exposed to.

Beyond terminology, careful thought needs to be given to the paradigm to use when presenting identity to users, as identities and services are abstract concepts that some users may not find easily understandable. Implementers may wish to keep such abstract concepts despite this, or may wish to examine attempts to map to real world paradigms, e.g. the idea of using "Identity Cards" that are held in the user's "Wallet", as used by the now defunct Microsoft Cardspace ([MS-CS]).

5.2. Services

Terminology around services is likely to be less of a problem than identity, but it will actually depend on what the service is. For example, each service could be simply described as "server", "system", etc. But for simplicity just the word "service" will probably usually suffice.

5.3. Identity to Service Mapping

The basic functionality of the Identity Selector is to create the correct combination of Identity and Service, so that the correct identity is chosen to create the credential for the GSS-EAP connection with the given service. Mapping is the process of creating this relationship between identity and service.

Depending on your perspective either each identity may be mapped to multiple services, or each service has multiple identities mapped to it. Thus any UI could present either perspective, or both.

6. Considerations around Management of Identities

One of the core features of an identity selector is the management of a user's identities. This section first looks at what information associated with an identity will need to manage, and then looks in detail at various usability considerations of this area.

6.1. Information associated with each Identity

The bare minimum set of information that **MUST** be stored about each identity to allow ABFAB authentication to take place is a single item:

- o NAI: The user's Network Access Identifier (see [RFC4282]) for this particular credential. For example, "joe@example.com". Note that the identity selector **MUST NOT** store different identities that use the same NAI. This is required as the NAI is the unique key that is used by the identity selector when interacting with the GSS-API mechanism for various reasons, for example, to allow the GSS-API mechanism to report back error or success statuses or to allow the application to request the use of a specific identity.

Next up is a small set of information that **SHOULD** be stored about each identity to allow the user to effectively select a particular identity:

- o Identity provider realm: The ABFAB realm of the identity provider. This is used as a key to look up the identity provider from the identity selector's list of identity providers, in order to access the trust anchor during verification of the identity provider.
- o Credential: Whatever is used by the users to authenticate themselves with a particular NAI. What exactly this will be will be dependent on the EAP method being used, but is likely to be something like a password or a certificate. Note that the identity selector **SHOULD** allow a user to store the credential. However, there are use cases where a user may specifically opt for this not to be "remembered", so the identity selector **MUST NOT** store the credential without confirmation from the user.

Finally, there is a set of optional information that **MAY** be stored about each identity that represent useful information for the user to have and could make an identity selector more usable. Note that this list is neither intended to be exhaustive or even particularly correct; any implementer is free to use whatever make sense in their implementation and conforms to good HCI/UX guidelines. Instead, it is simply a suggested starting point.

- o Friendly name for identity: To allow the user to differentiate between the set of identities represented in the Identity Selector. This should be editable by the user. The only restriction on this name is that it MUST be unique within that particular user's set of identities. For example: "Student username", "Google Account", "Work Login", etc.
- o Friendly icon for identity: To allow the user to differentiate between the set of identities they have they should be able to set an icon for that particular identity.
- o Password changing URL: The URL the user should visit should they need to change their password for this particular identity. For example, "http://www.example.com/passwordreset?identifier=myId".
- o Helpdesk URL: The URL this particular identity should visit to get contact details for the helpdesk of the organisation that issued this particular identity for when the user encounters issues and needs help. For example, https://www.example.com/helpdesk?identifier=myId.

6.2. Information associated with each Identity Provider

Identity providers are entities that may be shared across multiple identities. For instance, a person at a university may have one identity as a student and another identity as an employee, but a single identity provider makes access management decisions about both. In these cases, the identity selector MUST consider it an error if the trust anchor for the identity provider is different between the various identities managed by the single identity provider.

The bare minimum set of information that MUST be stored about each identity provider is:

- o Realm: The realm of the identity provider. This will uniquely identify the identity realm.
- o Trust anchor: For the identity selector to be able to verify that the Identity Provider it is going to talk to and attempt to authenticate against is the Identity Provider that it is expecting, and that it is not being spoofed in some way. This is likely to be an X.509 certificate [RFC5280], or a tuple of (trusted root certificate, servername in Subject or subjectAltName). Storing a credential without a relevant trust anchor allows for the possibility of a malicious attacker intercepting traffic and masquerading as the Identity Provider in question.

Identity providers also have a set of optional information that MAY be stored about each identity provider. This set includes, but is not limited to:

- o Friendly name for the identity provider: To allow the user to differentiate between the set of identity providers represented in the Identity Selector. This should be editable by the user. The only restriction on this name is that it MUST be unique within that particular user's set of identity providers. For example: "My University", "Google", etc.
- o Friendly icon for the identity provider: To allow the user to differentiate between the set of identity providers they have they should be able to set an icon for that particular identity provider.
- o Password changing URL: The URL the user should visit should they need to change passwords for identities in this realm. For example, "http://www.example.com/passwordreset".
- o Helpdesk URL: The URL the user should visit to get contact details for the helpdesk of the organisation that issued this particular identity for when the user encounters issues and needs help. For example, <https://www.example.com/helpdesk>.

Note that the password changing URL and helpdesk URL somewhat mirror the definitions of the same fields in the identity. The distinction is that the URLs in the identity SHOULD apply to the individual identity, whereas the URLs in the identity provider SHOULD apply to all identities that the identity provider defines. For example, an identity password change URL would provide a personalized experience of changing the password for the given identity, but the identity provider password change URL would direct the user to a page where the user would need to enter the individual identity that needs a new password.

If the identity contains no password change URL or helpdesk URL, the identity selector MAY present any corresponding URL from the identity selector instead. However, if the identity contains the URL, the identity selector SHOULD present the URL from the identity.

6.3. Storage of Identity Information

Since some of the information that makes up the identity is sensitive in nature (e.g. containing passwords), then this information SHOULD be stored and accessed securely. This might involve ensuring the credential information is held in encrypted form on device and accessed using a passphrase. For deeper integration into the system,

this could be done by using existing secure storage on the system such as Keychain on a Mac, the GNOME keyring on a GNOME based Linux device, or the Credentials Manager on Windows.

6.4. Adding/Association of an Identity

Users will have one or more identities given to them by organisations that they have a relationship with. One of the core tasks of an identity selector will be to learn about these identities and their identity providers in order to use them when it comes to authenticating to services on behalf of the user. Adding these identities could be done in one of three ways: manual addition, automated addition that is manually triggered, or automated addition that is automatically triggered. Each of these are discussed in more detail next.

Note that the term "association" or "addition" of an identity is used rather than "provisioning" of an identity, because while we actually are provisioning identities into the UI, provisioning is an overloaded term in the identity and access management space and could easily be confused with identity provisioning in the sense of the creation of the identity by the home organisation's identity management procedures.

6.4.1. Identity Provider Addition

6.4.1.1. Manual Identity Provider Addition

Allowing users to add an identity provider manually is technically the easiest method to get this information, but it is a method that has the greatest usability drawbacks - including some that create potential security issues. Most of the information required is relatively technical and finding some way of explaining what each field is to a non-technical audience is challenging (to say the least). This especially is the case for trust anchor information. Thus this method should be considered as a power-user option only, or as a fall-back should the other methods not be applicable. Implementers may well decide not to offer the manual option due to these drawbacks.

When this method is used, careful consideration should be given to the UI presented to the user. The UI will have to ask for all of the information detailed in Section 6.2.

Trust anchors present a particularly onerous challenge for the user to enter. For this reason, many identity selectors will want to implement a leap-of-faith acquisition of the trust anchor. For these leap of faith acquisitions, the identity selector SHOULD present the

user with the name of the realm that the identity selector is attempting to reach, the subject of the trust anchor certificate, details of the certification chain, and a fingerprint of the certificate. If the realm does not match the subject of the certificate, the identity selector MUST inform the user of the discrepancy. The identity selector MAY reject the leap-of-faith on its own, or MAY allow the user to proceed anyway. If the user proceeds anyway, the identity selector SHOULD urge the user to reject the leap-of-faith.

The area of verification of trust anchors is very important. An Identity Selector that allows for manual addition of identity provider information SHOULD try to ensure that trust anchor information is gathered and checked in a secure a manner as possible - where users have to enter and confirm all trust anchor information, or be required to explicitly agree to an insecure configuration if this is not done properly.

6.4.1.2. Manually Triggered Automated Identity Provider Addition

One way to bypass the need for manual addition of an identity provider - and all of the usability and security issues inherent with that approach - is to provide some sort of manually triggered, but automated, addition process. One approach to accomplishing this, for example, could be for an organisation to have a section on their website where their users could visit and be given piece of data that contains much or all of the relevant identity provider information for importing into the identity selector.

Additionally, the user SHOULD be given the opportunity to:

- o Supply or change the default friendly name and friendly icon for that identity provider - to allow the user to customise the identifier they use for that identity provider;
- o Reject the addition of the identity provider completely - to allow the user to back out of the association process in an intuitive way.

In this case, trust anchors would be directly provided through the automated addition process to help establish the trust relationship in a secure manner.

6.4.1.3. Fully Automated Identity Provider Addition

Many organisations manage the machines of their users using enterprise management tools. Such organisations may wish to be able to automatically add a particular user's identity provider to the

identity selector on their machine/network account so that the user has to do nothing.

This represents the best usability for the user - who wouldn't actually have to do anything. However, it can only work on machines centrally managed by the organisation.

6.4.2. Identity Addition

6.4.2.1. Manual Identity Addition

Allowing users to add an identity manually is relatively easy in comparison to adding an identity provider manually. If the identity provider is already known in the identity selector, then the identity selector can construct the NAI from the identity provider and a username. Thus the manual addition of an identity in a known realm needs to prompt the user only to pick the realm, to enter the username, and to enter the credential. If the identity provider is not known to the identity selector, the identity selector will provide the user with a way to define a new one as part of the identity addition.

There are two points at which a user could manually add an identity:

- o Asynchronously: the user could be allowed to, at any time, trigger a workflow of manually adding an identity. This represents the most flexible way of adding an identity since a user can perform this at any time. It does, however, also have inherent issues when it comes to verifying the newly added identity - see Section 6.6.
- o Just In Time: when connecting to a service which has no mapping to an existing identity, the user could be given an option to add a new one, as well as associating with an existing one. This seems to present a better user experience when it comes to verifying the newly added identity (see Section 6.6), however, it represents a less direct method of adding an identity. Users who have not yet added the appropriate identity to their identity selector may find it difficult to understand that they must try to access a particular service in order to add an identity.

Of course, implementers could support both styles of identity addition to gain the benefits of both and give flexibility to the user.

6.4.2.2. Manually Triggered Automated Identity Addition

Much like in the case of the manually triggered automated identity provider addition Section 6.4.1.2, an identity could be added to the identity selector through a user-initiated mechanism. To follow the example in the identity provider section above, the organization could enhance the identity provider addition web service to prompt for the user part of the NAI. The web service could then generate all of the data needed for adding both the identity provider and the identity.

It is reasonable to assume that any such automated addition service is likely to be organisation specific, so that the Issuing Organisation and realm part of the NAI will be constant, as would be the trust anchor information. The user part of their NAI will have been input on the web service. The password could be provided as a part of the provided data or the identity selector could prompt the user to enter it.

If the identity provider data contained in this identity to be added conflicts with an existing identity provider known to the identity selector, the identity selector SHOULD present the discrepancy to the user. The identity selector MAY reject the identity provider and identity on its own, or MAY allow the user to proceed anyway. If the identity selector allows the user to proceed anyway, the identity selector SHOULD urge the user to reject the leap-of-faith, and require the user to confirm the intent to proceed before proceeding.

Additionally, the user SHOULD be given the opportunity to:

- o Supply or change the default friendly name for that identity - to allow the user to customise the identifier they use for that identity;
- o Indicate whether or not the identity selector should always ask before using services with this identity - to customise the way in which the identity selector interacts with the user with this particular identity;
- o Reject the addition of the identity completely - to allow the user to back out of the association process in an intuitive way.

6.4.2.3. Fully Automated Identity Addition

Section Section 6.4.1.3 introduced the concept of using enterprise management tools to add an identity provider to the identity selector. These enterprise management tools could be used to add an identity that uses the identity provider added in the above manner.

The user would not need to decipher difficult to understand data entry screens.

However, having an identity automatically provided, including its password, does have some particular usability issues. Users are used to having to provide their username and password to access remote services. When attempting to access services, authenticating to them completely transparently to the user could represent a source of confusion. User training within an organisation to explain that automated population of their identity has been enabled is the only way to counter this.

6.5. Modifying Identity Information

This process is conceptually fairly similar to adding an identity, and thus shares many of the usability issues with that process. Some particular things are discussed here.

6.5.1. Manual Modification

An identity selector may allow a user to manually modify some or all of the information associated with each identity. The obvious items that SHOULD be allowed to be changed by the user are the friendly name, the friendly icon, and the credential, or password, associated with the identity.

The identity selector should restrict other modification of the information:

- o Identity Selectors SHOULD NOT allow the editing of the NAI of an identity or the trust anchor of an identity provider for items that have been added through automated means (Section 6.4.1.2, Section 6.4.1.3, Section 6.4.2.2 and Section 6.4.2.3).
- o Identity Selectors MAY allow the update of the trust anchor of identity providers that have stored the trust anchor through just in time manual addition, using another just in time retrieval of the trust anchor. Any identity selector that allows this update MUST inform the user of the change in the trust anchor, and advise the user that any unexpected change should be assumed to be an attack.
- o Identity Selectors SHOULD NOT allow manual modification of the password changing URL.
- o Identity Selectors SHOULD NOT allow manual modification of the helpdesk URL.

6.5.2. Automated Modification

To ease usability, organisations may wish to automatically provide updates to identity provider or identity information. For example, if the user's password changes it could automatically update the password for the identity in the user's identity selector, or if the trust anchor information changes (e.g. if a certificate is changed) it could be automatically pushed out to all users.

6.6. Verifying an identity

An inherent by-product of the ABFAB architecture is that an identity cannot be verified during the addition process; it can only be verified while it is in use with a real service. This represents a definite usability issue no matter which method of identity addition is used (see Section 6.4):

- o If the user has manually added the identity (see Section 6.4) they may have gone through the whole manual process with no errors and so believe the identity has been set up correctly. However, when they attempt to access a service, they may be given an error message, thus causing some amount of confusion.
- o If the user has had the identity populated into their identity selector, then there is a much greater chance of the identity information being correct. However, if any of the information is not correct, then there is the potential for confusion as the user did not add the information in the first place.

Also, if the identity information is incorrect the user may not know where the error lies, and the error messages provided by the process may not be helpful enough to indicate the error and how to fix it (see Section 8).

6.7. Removing an Identity

This is fairly similar to adding or modifying an identity, and thus shares many of the usability issues with those processes. Some particular things are discussed here.

6.7.1. Manual Removal

Allowing the user to manually delete an identity is probably the best way to achieve the goal. Any UI should allow for this option.

6.7.2. Automated Removal

While automated removal of an identity is a way of achieving the goal without having to interact with the user, the consequence is that things may disappear from the user's identity selector without them realising.

6.8. Storing an Identity with or without credentials

Sometimes, a user may wish to have the identity they wish to use with a service stored by the identity selector, but not the credential (e.g. password) that goes along with that Identity. The consequence of this is that when a user attempts to authenticate to a service for which an identity, but no credential, is stored, then the user would need to be prompted to manually enter the credential.

7. Considerations around Management of Service to Identity Mappings

A service to identity mapping tells the identity selector which identity should be used for a particular service. There is potentially a many-to-many association between identities and services since a user may wish to use one of their identities for many services, or more than one identity for a single service (e.g. if they have multiple roles on that service).

This potentially complex many-to-many association between identities and services is not easily comprehended by the user, and allowing the user to both manipulate it and control can be challenging. These obstacles are especially common when errors occur after an association has been made. In this scenario it is important that an identity can be disassociated with a service.

To further complicate the picture, users may wish for:

1. The identity to service mapping to be stored along with the credential, i.e. the user should always be authenticated to a particular service with a particular identity with no prompting.
2. The identity to service mapping to be stored but not the credential, i.e. the user should not be prompted to choose the identity for a particular service, but should be prompted to enter their credential for that identity.
3. The identity to service mapping to not be stored, i.e. the user should be asked which identity to use every time they authenticate to a particular service.

7.1. Associating a Service with an Identity

There needs to be a way for the user to create the service to identity association. It is advisable that this link be made only after the identity in question has authenticated with the service without any error.

There are a few ways this association could happen.

7.1.1. User-driven Manual Association

There are two ways in which manual association of an identity to a service could happen:

1. The identity selector MAY allow the user to associate a particular service with a particular identity manually, using the identity selector before they first attempt to use the service. This method is inadvisable, however, because not only might the identity in question not yet have authenticated successfully, the user would also need to know all the required technical details of that service beforehand, such as its GSS Acceptor Name.
2. On encountering a service new to the identity selector, the identity selector SHOULD pop up a dialogue box to the user asking if they would like to use an existing identity for this service (and might also allow them to create a new identity and use that).

7.1.2. Automated Rules-based Association

It would be beneficial from a usability perspective to minimise - or avoid entirely - situations where the user has to pick an identity for a particular service. This could be accomplished by having rules to describe services and their mapping to identities. Such a rule could match, for example, a particular identity for all IMAP servers, or a particular identity for all services in a given service realm. These rules could be configured as a part of the automated identity addition process described in Section 6.4.2.2 or Section 6.4.2.3.

7.1.3. Association Conflicts

The presence of rules-based associations brings with it potential conflicts in the rules. A non-exhaustive list of conflicts includes:

- o One rule applies to all services of a particular type, while another rule applies to all services within a particular domain. For example, one rule applies identity A to all IMAP services,

while another rule applies identity B to all services in the example.com domain.

- o One rule originates from enterprise management tools as described in Section 6.4.2.3, and another rule originates from manual addition.
- o The user has associated an identity with a service upon encountering the service for the first time, and later creates a rule that matches all services within that service's realm.

Identity selectors MUST order the precedence of rules as follows:

1. Manually created rules matching specific services and realms
2. Enterprise created rules matching specific services and realms
3. Manually created rules matching any service in a single realm
4. Enterprise created rules matching any service in a single realm
5. Manually created rules matching a single service in any realm
6. Enterprise created rules matching a single service in any realm

Identity selectors SHOULD notify the user whenever a new rule will take precedence over an existing rule.

7.2. Disassociating a Service with an Identity

A user MUST be able to disassociate an identity with a service - that is, to be able to remove the mapping without having to remove the identity.

For serious authentication errors, the identity selector SHOULD prompt the user to choose whether to disassociate the identity from the service or retain the association. The prompt SHOULD explain the nature of the error.

When such a serious authentication error occurs and the identity is selected by a rules-based association (Section 7.1.2), any disassociation prompt MUST inform the user that the identity was selected by a rule. The prompt SHOULD allow the user to retain the association, or to disassociate the rule altogether. The prompt MAY include a third choice, to create an exception so that the rule does not apply to this specific service.

As of this writing, there are no authentication failures that should cause the disassociation of an identity from a service.

7.3. Listing Services and Identities

A service listing should be considered in the identity selector which is both searchable and editable by the user.

7.4. Showing the Service that is requesting Authentication

When a user is attempting to authenticate to a service for the first time, there should be some indication given to the user as to which service is requesting authentication. In many cases, the service may be obvious (where the user has started the process of attempting to authenticate to a particular service), but in other cases this may not be obvious (e.g. if an authentication attempt is triggered by a timer or a specific event), and for this scenario some indication as to the requesting service is necessary.

7.5. Showing the Identity currently in use

It would be beneficial if, when using a service, the identity currently in use could be made visible to the user while they are using a specific service. This allows the user to identify which identity is used with a particular service at a particular time (the user may have more than one identity that they could use with a particular service) - so that they can then disassociate the pairing. This is especially useful when the identity is selected without any user prompt, because of a previous association.

Implementing such a feature may be hard, however, due to the layered nature of the ABFAB transaction - the identity selector will certainly know when successful (or failed) authentications to a particular service have happened, but after that it typically plays no further part in the use of the service. Therefore, knowing that a particular service is still using a particular identity in order to indicate this to the user would be challenging.

One approach that could be used would be to display OS notifications when an identity is used. The notification could include information such as the application requesting the identity, the service receiving the identity, and the identity used. Another approach could be for the identity selector to maintain a history of identity use.

7.6. Multiple Identities for a Particular Service

An Identity Selector should be able to deal with the case where a user has multiple identities associated with a single service. For example, upon receiving a request for authentication to a service that multiple identities are configured for, ask the user which of the identities should be used in this instance.

7.7. Not using ABFAB for a Particular Service

There may be cases where a user does not wish to use ABFAB based authentication at all to a particular service, even though it is ABFAB enabled. To support this, the identity selector would have to allow the user to choose not to use ABFAB when they attempt to authenticate to a service. It would be desirable if the user could also flag that this should be remembered.

8. Handling of Errors

Errors during the ABFAB authentication process can happen at any of the many layers - they could be GSS-API errors, EAP errors, RADIUS/RadSec errors, SAML errors, application errors, etc. ABFAB based technologies are limited in error handling by the limitations in the protocols used.

8.1. Errors in GSS-API

All GSS-API calls are necessarily instantiated from within the calling application. For this reason, when an error occurs the error is passed back to the application in order for it to deal with it. To retry, the application needs to re-initiate the GSS-API call. Unless the application has been written to deal with this properly, this process can be very tedious for a user and cause them opt out of what they are trying to accomplish. In addition to this, the application may not display the error messages to the user. Even when the application does display the errors, the messages themselves may not be useful enough for the user to decipher what has gone wrong.

Two extensions to GSS-API are suggested for the consideration of the kitten working group:

- o GSS-API should provide a method for applications to invoke to indicate that the application has displayed the last error to the user.

- o GSS-API should provide a method for applications to invoke to indicate that the authentication succeeded, but is insufficient for the task at hand and needs to be retried.

8.1.1. Log of Errors

The Identity Selector can improve the general GSS-API error reporting experience by displaying a list of errors experienced by ABFAB applications. When an application error occurs, the EAP mechanism MAY record that error. If the mechanism records these errors, the Identity Selector MAY display these errors to the user. Thus, the user will have a single place to go to view all of the errors that a user experiences across all applications. Therefore, an Identity Selector that implements an error display SHOULD present the user with the context of the error, including the calling application and the time.

8.2. Examples of errors

To give an idea of the range of errors that might be seen, consider the following non-exhaustive set of potential errors.

Identity Association/Verification Errors:

- o The credentials presented to the IdP were not able to be verified - e.g. wrong username/password.
- o The Trust Anchor for the IdP was invalid.

Service Errors:

- o The IdP recognizes the client, but decides not to authorize it for this service.
- o The EAP session succeeds, but the RADIUS system sends access-reject to the Relying Party
- o The RADIUS system succeeds, but the Relying Party rejects the session. For instance, the SAML part of the session could contain an error that causes the Relying Party to reject the client.
- o The Identity might have been successfully authenticated, but the user might not have authorisation to use the service or privilege levels within the service they are attempting to use. For instance, the Identity could authorise the use of an operating system as an unprivileged user, which would prevent the user's goal of managing the hard drives.

Other Errors:

- o The IdP didn't respond to the Service.
- o The IdP didn't respond to the Client.
- o Network errors.
- o Timing errors.

9. Handling of Successes

It is of course hoped that the identity selector will have to occasionally handle successes as well as errors. This section has some brief discussion about some areas you might want to think about.

9.1. Reporting Authentication Success on First Use of Identity

The first time an identity is used with a service, it would be good practice to visually indicate in some way that the process has been successful, in order that the user understands what is happening and is then prepared for future authentication attempts.

9.2. Reporting Authentication Success

On an on-going basis you may or may not wish to indicate visually to the user a successful authentication to a service. This relates to Section 7.5.

10. Other Considerations

This section briefly discusses other considerations that you might want to think about that don't fit in any of the other categories.

10.1. Identity Selector Taking Focus

When an ABFAB authentication request is triggered, and where it needs input from the user, the Identity Selector should take focus in some way so that it is clear to the user that they need to do something to proceed.

10.2. Import/Export of Credentials

For various reasons, an identity selector implementation might want to include functionality that allows for the export/import of identities and service to identity mappings. This could be for backup purposes, to allow a degree of mobility between identity selector instances, etc.

If providing this functionality, it would be advisable that the credential store that is the result of the export should be secure - encrypted and password protected - given the nature of the information.

11. Security Considerations

Most security considerations are ones relevant to the use of GSS-EAP and are detailed in [I-D.ietf-abfab-arch]. There are, however, a few specific sets of security considerations related to the UI implementation.

First, as discussed earlier, the Identity Selector should use a Trust Anchor to authenticate the IdP before it sends the users credentials to it. Having no Trust Anchor information at all, or an incorrect Trust Anchor, can enable the possibility of someone spoofing the IdP and harvesting credentials sent to it. So, how this Trust Anchor is configured and managed can have major security implications:

- o The most secure way for a Trust Anchor to be configured is to have it provisioned alongside the other identity information in an enterprise provisioning scenario. This allows for the correct Trust Anchor to be configured with no user input required. However, thought needs to be given to Trust Anchor expiry and consequent requirement for regular reprovisioning of identity information.
- o Another way that is potentially secure would be to allow the user to discover the Trust Anchor information out of band and manually input this information into the Identity Selector. This is only secure, however, for those users who understand what they're doing in this scenario; pragmatically, this is unlikely to be the case for many users so is not a recommended approach for the average user.
- o A pragmatic approach would be leap of faith, whereby no Trust Anchor information is initially provisioned, and the first time the Identity Selector connects to the IdP it remembers the Trust Anchor information for future use. This doesn't mitigate against spoofing of an IdP in the first instance, but would enable mitigation against it for all future connections.
- o Finally, there may be interesting ways to leverage technologies such as DANE [RFC6698] to store the Trust Anchor for an IdP in DNS.

Secondly, the storage of the user's credentials by the Identity Selector should be done in a secure manner to mitigate against people

taking unauthorised control of the device being able to gather these credentials. Use of a secure credential storage mechanism, such as the GNOME Keyring on Linux, or Keychain on the Mac, are recommended.

12. Privacy Considerations

Since the ABFAB system facilitates the sharing of identifying information about a user, the undesired sharing of information is a real concern. Most of the privacy considerations lie outside the scope of the Identity Selector UI, which neither controls nor sees which attributes of an identity will be shared with a service. In essence, the only control that the Identity Selector has is whether or not a given identity will be shared with the service.

However, the selection of identity does warrant privacy considerations. Any automated choice of identity for a service will share information, potentially inappropriately. Examples of this include:

- o Rules that apply to a service across all realms will cause an identity choice, even for realms the user would actually prefer to avoid interacting with at all.
- o Storing a default for a particular service and realm will cause the identity to be selected in that situation going forward, even if the situation or application does not warrant that. For instance, a web browser in privacy mode ideally should not know of any saved identity association choices.

Even appropriate choices of sharing an identity with a service leaks information about the user. The desired service and the identity provider must communicate with each other to perform an authentication. Even if the authentication fails, the service will know the realm of the user credential, and the Identity Provider will know the realm, and maybe the service, that the user tried to access. For services with fallback authentication mechanisms, the system may try and fail to authenticate the user, thus sharing the realm information noted above, without the user being aware this has happened.

13. IANA Considerations

This document does not require actions by IANA.

14. Contributors

The following individuals made important contributions to the text of this document: Sam Hartman (Painless Security LLC), and Maria Turk (Codethink Ltd).

15. Acknowledgements

Thanks to Jim Schaad, Stefan Winter, David Chadwick, Kevin Wasserman, Alejandro Perez-Mendez, Ken Klingenstein, and Dave Crocker for feedback and suggestions.

16. References

16.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4282] Aboba, B., Beadles, M., Arkko, J., and P. Eronen, "The Network Access Identifier", RFC 4282, DOI 10.17487/RFC4282, December 2005, <<http://www.rfc-editor.org/info/rfc4282>>.
- [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, DOI 10.17487/RFC5280, May 2008, <<http://www.rfc-editor.org/info/rfc5280>>.
- [RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", RFC 6698, DOI 10.17487/RFC6698, August 2012, <<http://www.rfc-editor.org/info/rfc6698>>.
- [I-D.ietf-abfab-arch] Howlett, J., Hartman, S., Tschofenig, H., Lear, E., and J. Schaad, "Application Bridging for Federated Access Beyond Web (ABFAB) Architecture", draft-ietf-abfab-arch-12 (work in progress), February 2014.
- [I-D.ietf-abfab-usecases] Smith, R., "Application Bridging for Federated Access Beyond web (ABFAB) Use Cases", draft-ietf-abfab-usecases-05 (work in progress), September 2012.

16.2. Informative References

- [MS-CS] Brown, K., "The InfoCard Identity Revolution", July 2006, <<https://technet.microsoft.com/en-us/magazine/2006.07.infocard.aspx>>.

Appendix A. Change Log

Note to RFC Editor: if this document does not obsolete an existing RFC, please remove this appendix before publication as an RFC.

IETF draft -03 to ietf draft -04

1. Document service errors.
2. Document GSS error handling, including a request for a couple of new GSS methods, and maintaining a log of all GSS errors for later viewing.

IETF draft -02 to ietf draft -03

1. Tidying up language throughout.
2. Added the idea of an identity provider object within the identity selector, and moved the trust anchor property from the identity to the identity provider.
3. Added restrictions on manual modification of automatically added identities and identity providers.
4. Added precedence between identity association rules.
5. Incorporated many comments from the mailing list.
6. Added privacy considerations section.

IETF draft -01 to ietf draft -02

1. Tidying up language throughout.
2. Finished remaining TODOs - largely in the error handling section.
3. Added security considerations section.

IETF draft -00 to ietf draft -01

1. Tidying up language throughout
2. Doing some of the TODOs
3. Added language that tries to explain that this document is not a substitute for good HCI/UX design.

4. Changed terminology slightly to avoid confusion between an identity selector "mechanism" and a GSS-API mechanism.
5. Added a caveat about the potential for the UI to show the identity currently in use for a particular service.
6. Added a requirement that the identity selector must not store the same NAI for multiple identities.
7. Stopped talking about "provisioning" after saying that I wouldn't talk about "provisioning".

Draft -04 to ietf draft -00

1. Adding brief discussion of identities vs identifiers (Ken).
2. Changing assumption about credentials having a password in favour of more generic text for other auth types.
3. Adding discussion of storage of identity information.
4. Added sections on dealing with multiple identities per service, remembering credentials, remembering not to use ABFAB.
5. Added small section on ID selector needing to take focus in some way.

Draft -03 to draft -04

1. Addressing various comments from Jim and Stefan.

Draft -02 to draft -03

1. Bumping version to keep it alive.

Draft -01 to draft -02

1. Completed the major consideration sections, lots of rewording throughout.

Draft -00 to draft -01

1. None, republishing to refresh the document. Other than adding this comment...

Appendix B. Open Issues

Note to RFC Editor: please remove this appendix before publication as an RFC.

Authors' Addresses

Dr. Rhys Smith
Cardiff University
39-41 Park Place
Cardiff CF10 3BB
United Kingdom

Phone: +44 29 2087 0126
EMail: smith@cardiff.ac.uk

Mark Donnelly
Painless Security
14 Summer Street
Suite 202
Malden, Massachusetts 02176
United States

EMail: mark@painless-security.com

ABFAB
Internet-Draft
Intended status: Informational
Expires: September 7, 2014

L. Nordberg
NORDUnet
J. Howlett
JANET(UK)
March 06, 2014

Ephemeral keying for ABFAB
draft-linus-abfab-ephemeral-keying-01

Abstract

This document describes how EAP-GSS provides forward secrecy by encrypting each session in an ephemeral key generated in the initial state of the context establishment. This Diffie-Hellman key is shared by the initiator (EAP peer) and acceptor (EAP authenticator).

The goal is to stop a passive attacker with access to the traffic between an ABFAB user and the service she uses (Relying Party), from getting access to key material and information linkable to the user or from being able to fingerprint the user.

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 September 7, 2014.

Copyright Notice

Copyright (c) 2014 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
2. Information potentially accessible to a passive observer . .	2
2.1. RADIUS	3
2.2. EAP method	3
2.3. GSS-API data	3
3. Solution	4
3.1. Why do this at the GSS-EAP layer	4
4. Keying algorithm	5
5. Costs	5
6. Open questions	5
7. Security considerations	5
8. IANA considerations	5
9. Contributors	5
10. Normative References	6
Authors' Addresses	6

1. Introduction

The ABFAB architecture [I-D.ietf-abfab-arch] defines a GSS-API mechanism for the Extensible Authentication Protocol [RFC7055]. This mechanism provides support for the security services offered by the GSS-API, including the confidentiality of context tokens. This confidentiality service is available once a GSS context has been negotiated successfully between the initiator and acceptor.

However there is a possibility that a passive observer could extract information from this negotiation that could potentially compromise aspects of the confidentiality of the context tokens and/or the privacy of the initiator and/or acceptor.

This document defines an extension to [RFC7055] to deny a passive observer access to this information by encrypting the tokens used to establish the GSS context.

2. Information potentially accessible to a passive observer

This section describes the information available to a passive observer of an [I-D.ietf-abfab-arch] authentication, working from the lowest layers of the protocol stack upwards.

2.1. RADIUS

The realm component of the NAI [RFC4282] is generally exposed. While the user name component of the NAI is easily anonymised, the realm (which effectively names the user's identity provider (IdP)) will provide a strong indication of the organisational affiliation of a user.

In the event that RADIUS/UDP is being used instead of RADIUS/TLS, not only do the intermediate proxies between the acceptor and the IdP have access to the EAP MSK but a passive observer does too. Knowledge of the MSK could facilitate the compromise of the GSS context, which is derived from this key, potentially allowing decryption of the GSS session.

2.2. EAP method

The EAP methods most commonly used with [RFC7055] use X.509 server certificates to authenticate the IdP. This certificate will include information identifying the IdP's server.

A passive observer may also be able to fingerprint the EAP implementation [FIXME].

In cases where a TLS-based EAP method is used, a passive observer may be able to fingerprint the client based on TLS session resumption, for example as described in [RFC5077] section 5.8.

2.3. GSS-API data

A variety of information is available at the GSS-API layer.

- o The acceptor name is carried in name requests and responses during the initial phase. This can be used for fingerprinting users since it indicates what service is requested and supplied. In settings where the endpoint's IP addresses and other identifying information don't link the user to the service, exposing the acceptor name is detrimental to privacy.
- o GSS channel bindings are also available in the extensions state; these bindings typically identify the acceptor to the initiator.
- o The currently defined flags leak information about which application protocol is being used and pose a threat to user privacy. Future flags might increase this threat.
- o Finally the mechanism MIC is also exposed and error subtokens are also exposed [FIXME].

3. Solution

Generate a Diffie-Hellman key in the initial state of the context establishment and use it to encrypt other context tokens. Note that the DH key, shared by initiator and acceptor, is unique per GSS-API session, not per context token. [Elaborate on why?]

[describe where in initial the DH key exchange happens and how; point at general description? copy from existing standard?]

[describe how we signal algorithm and key size]

[describe the use of a nonce/sequence number for temporality, either in the key or in the payload, covered by the MIC and verified by the other end - mitigates replay, reflection and reordering attacks]

[describe how we derive a symmetric key from the DH key and encrypt the context token (perhaps in a GSS "wrap token"?)]

[describe how to mix in the DH key with the MSK to form the CRK (7055 sect 6) - this will make a MITM kexing with both ends unable to create a MIC which validates properly (and a MITM relaying DH kex will not know the key and thus not the CRK)]

3.1. Why do this at the GSS-EAP layer

Using a short lived key for providing confidentiality between an ABFAB client and the IdP could arguably be done at the EAP layer rather than at the GSS-API layer. A general solution for EAP would give better protocol reuse.

EAP methods run between the EAP peer and server. A Diffie-Hellman key exchange between these endpoints can not start with the first message sent from the client since the client doesn't talk to the EAP server (the IdP) directly and can not be helped with doing that until the EAP authenticator knows where the IdP is to be found. Most of the mentioned leaks at the GSS-API layer would thus still be present in this solution.

[maybe expand on how TEAP [draft-ietf-emu-eap-tunnel-method] could solve the problem of AAA proxies learning the MSK, impersonating the RP]

An alternative place to protect ABFAB authentication with a short lived key would be in the application level protocol. While some applications are using protocols already able to protect the GSS-API traffic using a TLS session with an ephemeral key (XMPP, IMAP, SMTP) it's not mandatory to use such a tunnel. Other applications use protocols which might be hard to protect in a tunnel (NFS, SSH).

4. Keying algorithm

This section defines an algorithm, based on the Diffie-Hellman protocol, enabling the initiator and acceptor to negotiate a shared key during the initial phase of the GSS context establishment. This key is used to encrypt all subsequent context tokens. The key is unique per GSS-API session, and is not rotated for each successive context token. [Elaborate on why not?]

5. Costs

- o This will cost FIXME extra round trips.
- o [No new GSS mech. Thus no complexity cost of picking the right one.]

6. Open questions

- o Should we make the ephemeral keying and encryption optional?

Might have to - asking the list about breaking backward compatibility.
- o Bid down attacks - detect, prevent

Fascinating idea from Sam: 6067 CB implementing 5056 CB could detect MITM before end of extension state (MIC).
- o Include the nonce/sequence number in tokens or fold it into the key?

7. Security considerations

TBD

8. IANA considerations

TBD.

9. Contributors

The whole idea of adding ephemeral keys to ABFAB was suggested by Sam Hartman who also contributed substantial ideas and discussions on this subject.

Jim Schaad has made several valuable comments with corrections and suggestions.

10. Normative References

- [I-D.ietf-abfab-arch]
Howlett, J., Hartman, S., Tschofenig, H., Lear, E., and J. Schaad, "Application Bridging for Federated Access Beyond Web (ABFAB) Architecture", draft-ietf-abfab-arch-12 (work in progress), February 2014.
- [RFC2743] Linn, J., "Generic Security Service Application Program Interface Version 2, Update 1", RFC 2743, January 2000.
- [RFC3748] Aboba, B., Blunk, L., Vollbrecht, J., Carlson, J., and H. Levkowetz, "Extensible Authentication Protocol (EAP)", RFC 3748, June 2004.
- [RFC4282] Aboba, B., Beadles, M., Arkko, J., and P. Eronen, "The Network Access Identifier", RFC 4282, December 2005.
- [RFC5077] Salowey, J., Zhou, H., Eronen, P., and H. Tschofenig, "Transport Layer Security (TLS) Session Resumption without Server-Side State", RFC 5077, January 2008.
- [RFC7055] Hartman, S. and J. Howlett, "A GSS-API Mechanism for the Extensible Authentication Protocol", RFC 7055, December 2013.

Authors' Addresses

Linus Nordberg
NORDUnet

Email: linus@nordu.net

Josh Howlett
JANET(UK)

Email: Josh.Howlett@ja.net