Best practices and recommendations for attribute translation from federated authentication to X.509 credentials

Published Date:	13-06-2017
Revision:	1.0
Work Package:	JRA1
Document Code:	AARC-JRA1.4l
Document URL:	https://aarc-project.eu/wp-content/uploads/2017/03/AARC-JRA1.4l.pdf

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

Table of Contents

1 Introduction 3

2 Authentication and authorization information separation 3

3 Translating authentication information 4

3.1 Description of End-Entity Certificates 4

3.2 Converting SAML attributes into a subject DN 5

3.2.1 Recommendations for defining the user CN RDN from IdP attributes 5

3.2.2 Recommendations for defining the O RDN from IdP attributes fields 6

4 Translating group information 7

5 References 8

6 Glossary 8

Table of Figures

Figure 1: Authentication and authorization information provision for X.509 credentials 4

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

1 Introduction

This document is a summary of the findings produced by the AARC project on the translation of authentication and authorization attributes from federated authentication credentials to X.509 credentials. This document shows examples of translation between SAML2 attributes and X.509, but the same principles are applicable to OpenID Connect sources of attributes.

The goal of this document is to suggest a common way to encode authentication and authorization in X.509 credentials, to increase the re-usability and interoperability of X.509 credentials generated by token translation services.

2 Authentication and authorization information separation

In the context of this document:

- Authentication information: The information that defines the identity of the user, such as, for example, name, unique identifier (UID) or organization.
- Authorization information: The information that is used to define whether a user is entitled to access a service: group membership (virtual organization membership), sub-group membership, or roles.

Although this may not be the actual implementation for every use case, we will assume that every authorization model within a community can be modelled as users' membership to groups created in the organization.

One of the main recommendations described in this document is that authorization and authentication information should be provided by separate entities, similar to what can happen in a federated identity scenario, where the IdPs provide the authentication information and the authorization attributes are provided by attribute authorities managed by the communities or other third parties. A similar approach should be implemented for the X.509 credentials, as shown in the following figure.

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

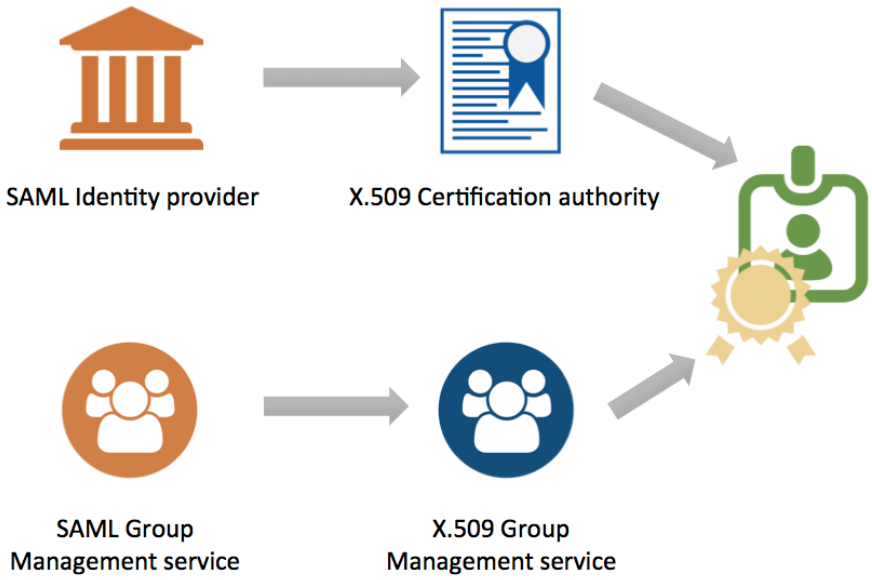

Figure 1: Authentication and authorization information provision for X.509 credentials

In a classic scenario, a Certification Authority (CA) plays the same role as a SAML or OpenID Connect Identity Provider (IdP); it creates a personal certificate with the authentication information of the user, while the group membership attributes are used by an X.509 group management tool to extend the user certificate with group membership information. For online CAs, the information from the IdPs is used by the CA to create the personal certificate. How to do that in a reliable and consistent way is the topic of this document.

3 Translating authentication information

IGTF [IGTF] is the most relevant X.509 certification authorities' federation for research infrastructures and e-infrastructures worldwide. The CAs federated in IGTF are commonly accepted by EGI, PRACE, EUDAT, OSG and many other service provider federations in the world. This section will focus on the IGTF profile for X.509 certificates as described by the existing policies, see [IGTF-PKI] and [GWD-R.225].

3.1 Description of End-Entity Certificates

The following are (some of) the attributes that are permitted in the end entity (EE) certificate DN:

DC: Domain

C: Country

ST: State or Province

L: Locality

O: Organization

OU: Organizational Unit

CN: Common name

For generic CAs a subject is uniquely identified by the combination of the subject DN and the issuer DN. Within the IGTF, this is further restricted and the subject DN itself is the unique non-reassigned identifier assigned to an end-entity (i.e. certificates with the same subject DN issued by different CAs MUST belong to the same end entity).

The following is an example of a complete certificate issuer and subject DN pair:

Issuer: /DC=org/DC=cilogon/C=US/O=CILogon/CN=CILogon Basic CA 1

Subject: /DC=org/DC=cilogon/O=Nikhef/CN=John Doe A12345

The issuer line defines the CA within the IGTF, while the subject line is a set of attributes specific of the user. In the remainder of this document we will only discuss the latter, as it is the relevant part for this document.

The subject DN starts with one or more Relative Distinguished Name (RDN) components uniquely (within IGTF) describing the organisation that operates the CA. The remaining RDNs MUST include a commonName (CN) attribute, which should be unique within the namespace, but typically an Organization (O) RDN describing the organization to which the user is affiliated, is also included there.

3.2 Converting SAML attributes into a subject DN

In order to produce suitable subject DNs for certificates issued by online CAs, where the subject DN must be formed from the available set of SAML attributes or corresponding OpenID Connect claims, certain guidelines must be followed. This section describes a reasonable and practical set of recommendations for doing this. The guidelines described here are mostly based on the practical experiences with the online CAs such as CILogon [CILOGON] and RCauth.eu [RCauth].

3.2.1 Recommendations for defining the user CN RDN from IdP attributes

1. The CN attribute shall preferably be derived from a combination of the SAML attributes (or their appropriate protocol equivalents) "displayName" and eduPersonUniqueId (ePUID)
    a. If eduPersonUniqueId is not released, eduPersonPrincipalName (ePPN) shall be used
    b. If eduPersonPrincipalName is also not released, eduPersonTargetedID (ePTID) or SAML NameID shall be used, on condition that it is *persistent* (as opposed to *transient*).

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

c. If displayName is not released, the value of givenName followed by a space and the value of sn shall be used. If those are also not released, the value of the commonName attribute shall be used.

How exactly the CN RDN is built up from these attributes is outside of the scope of this document, but the method or methods used must comply with these recommendations.

2. The CA must ensure that the CN is unique within their namespace.
3. RFC 5280 limits the length of the CN field to 64 characters. The string obtained as a combination of displayName+ePUID will typically be longer than this limit. A suitable solution for this problem is to re-hash the ePUID to a shorter, but globally unique version of the user's UID.
4. Since the CA relies on the received set of attributes, it must employ heuristics for guaranteeing that the same CN is only used for the same physical entity. This means that it could need to add further uniqueness guaranteeing components to the CN if it suspects, for example, non-uniqueness of the attributes under 1.
5. There are a few important practical consequences of the above points 3 and 4:
    a. It is difficult for the e-infrastructure to map reliably the user's UID provided by the IdP to the certificate DN. The CA therefore MUST provide this information as part of an incident response procedure.
    b. Even though point 4 might make it impossible to obtain the actual DN, the CA should publish publicly the re-hashing algorithm in order to aid the relying parties; any change to the algorithm should follow a clear procedure to communicate the changes to all stakeholders. In this way the e-infrastructure can calculate the re-hashed user ePUID used in the certificate and store the information internally, for a quicker incident response procedure.
    c. The full ePUID (or corresponding attributes) of the user could be stored in the certificate as a certificate extension (e.g. using otherName Subject Alternative Names) but this information will not appear in the subject DN. Since the subject DN is usually the only information about a user's certificate stored in the logs, e-infrastructures may want to autonomously map ePUID to the re-hashed user UID used by the CA, thereby improving traceability.

3.2.2 Recommendations for defining the O RDN from IdP attributes fields

The organization (O) RDN contains information about the organization of the user. The certification authority should add the value of one of the following as a separate organization RDN in the certificate DN, in order of preference:

1. schacHomeOrganisation attribute (urn:oid:1.3.6.1.4.1.25178.1.2.9)
2. organisationDisplayName in the IdP metadata
3. domain name of the IdP entity ID URL, or whole entityID in case of a URN

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

4 Translating group information

Group information should not be directly encoded in the user personal certificate. In the X.509 PKI certificate profile, additional attributes can be expressed through an Attribute Certificate (or Authorization Certificate, AC), which is signed by the attribute authority's EE certificate (usually the host certificate of the authority releasing the AC). The most common implementation of this is that used by VOMS, which is the only one we will consider here.

The advantages of keeping authentication and authorization information separated in two distinguished certificates are the following:

- The user certificate can be re-used, without changing the DN, forwith multiple VOs, in the not-unlikely scenario of users with multiple memberships (this includes also different roles or subgroups in the same VO).
- ACs can have a shorter lifetime than the EE certificate. Identity information changes less frequently than authorization information.

The user UID (the certificate DN) does not change depending on the VO used by the user.

ACs are usually embedded within a delegated proxy certificate, in the form of a special non-critical extension (with OID 1.3.6.1.4.1.8005.100.100.5, see [GFD-I.182]). In this way, the user's proxy certificate chain contains all the authentication and authorization information needed by the services in a single X.509 artefact.

Recommendations for the translation of authorization entitlements into X.509 Attribute Certificates

- Group membership information should be translated into a (VOMS) AC separate from the EE certificate. The AC should be included in a delegated proxy certificate as a non-critical extension. See [RFC3820], [RFC5755] and [GFD-I.182].
- The rules for translating these from SAML attributes, should follow the guidelines as set out in [AARC-JRA1.4A].
- One proxy certificate should contain information about a single VO to uniquely scope the group information, i.e. it should contain only a single AC.
- VOMS expresses groups, sub-groups and roles within the group in the form:

/<vo>/<group>/<sub-group>/../Role=<role>

- An AC should be short-lived, and its validity should not exceed the proxy certificate lifetime.

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

5 References

[CILOGON]	CILogon http://www.cilogon.org
[IGTF]	Interoperable Global Trust Federation https://www.igtf.net
[RCauth]	RCauth website http://rcauth.eu

6 Glossary

AC	Attribute/Authorization Certificate
CA	Certification Authority
eduPUID	eduPersonUniqueID
EE	End Entity
IdP	Identity Provider
IGTF	Interoperable Global Trust Federation
OGF	Open Grid Forum
OIDC	OpenID Connect
RFC	Request for Comments
SAML	Security Assertion Markup Language
SAML2	SAML Version 2.0
UID	Unique Identifier
VO	Virtual Organization

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).