Guidelines for expressing group membership and role information (AARC-G069)

Publication Date	2022-04-11
Authors:	Valeria Ardizzone, Dominik František Bučík, Marcus Hardt, Stefan Helmert, Jens Jensen, Ivan Kanakarakis, Christos Kanellopoulos, Nicolas Liampotis (ed.), Mikael Linden, Mischa Sallé
Document Code:	AARC-G069
DOI:	10.5281/zenodo.6533400
Community:	Architecture Area

Abstract

Information about the groups a user is a member of is commonly used by relying parties in order to authorise user access to protected resources. This document provides guidelines for expressing group membership and role information across AARC BPA-compliant AAI services. Specifically, it defines a URN namespace for expressing this information using common identity federation protocols, namely SAML and OpenID Connect/OAuth2.

Status of this document

This document provides guidelines for expressing group membership and role information and supersedes [AARC-G002]. The specification introduces the following changes compared to AARC-G002:

- Specifying the `<AUTHORITY>` component is optional; the use of `<AUTHORITY>` is deprecated and may be removed in a future revision.
- Group and role membership information is released through the `entitlements` claim [RFC9068] (instead of `eduperson_entitlement`) in the case of OpenID Connect relying parties.
- This specification adds rules for handling reserved or special characters (e.g. spaces, UTF-8 characters) when encoding URN values expressing group membership and role information.
- This specification adds case normalisation rules for encoding URN values expressing group membership and role information.
- This specification adds guidelines for the release of implied group membership information.

A list of all AARC community guidelines and the latest revision of each guideline can be found at https://aarc-community.org/guidelines.

Table of Contents

1 Introduction

Information about the groups a user is a member of is commonly used by Relying Parties (RPs) to authorise user access to protected resources. Apart from the group information that is managed by the user's home IdP, research communities usually operate their own group management services. Such services often act as Attribute Authorities, maintaining additional information about the users, including Virtual Organisation (VO) membership, group membership within VOs, as well as user roles. It is therefore necessary that all involved RPs and IdPs/AAs can interpret this information in a uniform way. Specifically, the following challenges need to be addressed:

- Standardising the way group membership information is expressed **syntactically**; for example, representing group membership as Uniform Resource Names (URNs) within a specific namespace and a set of rules for the Namespace Specific String (NSS) portion.
- Indicating the entity that is **authoritative** for each piece of group membership information
- Expressing **VO membership** and **role** information
- Supporting **group hierarchies** in group membership information
- Scoping: The rationale behind scoping is to prevent clashes between groups that are managed by different VOs/administrative domains. This eliminates the need for syntactical and semantic group information harmonisation among different communities. An added benefit is that scoping allows easy filtering of group values that can be used by RPs for quick authorisation decisions.

Harmonisation of naming for groups, hierarchy and use of ontologies within different scientific domains is explicitly excluded from this specification.

1.1 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].

2 General guidelines

This section describes a URN namespace for expressing group membership and role information in different federated identity protocols.

An attribute value expressing group membership and role information has the following syntax (components enclosed in square brackets are OPTIONAL):

```
<NAMESPACE>:group:<GROUP>[:<SUBGROUP>*][:role=<ROLE>][#<AUTHORITY>]
```

where:

- `<NAMESPACE>` is in the form of `urn:<NID>:<DELEGATED-NAMESPACE>[:<SUBNAMESPACE>*]`, where
  - `<NID>` is the namespace identifier associated with a URN namespace registered[1] with IANA, as per [RFC8141], ensuring global uniqueness. Implementers can and should use one of the existing registered URN namespaces, such as `urn:geant`[2] and `urn:mace`[3]
  - `<DELEGATED-NAMESPACE>` is a URN sub-namespace delegated from one of the IANA registered NIDs to an organisation representing the e-infrastructure, research infrastructure or research collaboration. It is recommended that a publicly accessible URN value registry for each delegated namespace is provided.

  A `<NAMESPACE>` can have a variable number of elements. For example `urn:geant:edugain`, `urn:geant:nikhef.nl` and `urn:geant:nikhef.nl:idm` are all valid `<NAMESPACE>` values.
- the literal string "`group`" indicates a value expressing group membership information;
- `<GROUP>` is the name of a Virtual Organisation (VO), research collaboration or a top level arbitrary group. Group names MUST be unique within a given namespace;
- an optional list of `<SUBGROUP>` components represents the hierarchy of subgroups in the `<GROUP>`;
- the optional `<ROLE>` component is scoped to the rightmost (sub)group; if no subgroup information is specified, the role applies to the top level group/VO;
- the `<AUTHORITY>` can be specified in the optional f-component[4] of the URN. Specifying the `<AUTHORITY>` is deprecated by this specification and may be removed in a future revision. When specified, the `<AUTHORITY>` MUST be a non-empty string introduced by the number sign ("#") character and terminated by the end of the URN. For example, it can be the FQDN of the group management system that is

[1] Generic top level namespaces require IANA approval as per Section 6.2 of [RFC8141]: https://www.iana.org/assignments/urn-namespaces/urn-namespaces.xhtml
[2] https://www.geant.org/Services/Trust_identity_and_security/Pages/NamespaceRegistry.aspx
[3] https://www.internet2.edu/products-services/trust-identity/mace-registries/urnmace-namespace/
[4] URN f-component: https://tools.ietf.org/html/rfc8141#section-2.3.3

responsible for the identified group membership information. As described in Section 2.2, the `<AUTHORITY>` MUST NOT be taken into account when determining equivalence[5] of URN-formatted values expressing group membership and role information.

Each group membership attribute value represents a particular position of the user within a VO, research collaboration or generally a top level arbitrary group. A user may be a member or hold more specific roles within the groups associated to this top level group. Groups are organised in a tree structure, meaning that a group may have subgroups, which in turn may have subgroups, etc.

This hierarchical structure implies that if someone is a member of a subgroup, then they are also a member of the parent group. For example:

`<NAMESPACE>:group:parent:child`

implies membership in `parent`, i.e.:

`<NAMESPACE>:group:parent`

A relying party receiving `<NAMESPACE>:group:parent:child:grandchild` MUST interpret this as membership of grandchild (a subgroup of child), membership of child (a subgroup of the parent), and membership of the parent. This type of group membership of child or parent is commonly termed implied, implicit or indirect group membership.

When an entity (e.g. SP-IDP-Proxy) needs to indicate membership of one or more (sub)groups in a group hierarchy, it MUST send assertions for each of the rightmost (sub)groups it wants to indicate. In addition, the entity MAY send any or all of the implied group memberships in that hierarchy explicitly. For example, it MAY send `<NAMESPACE>:group:parent` and `<NAMESPACE>:group:parent:child` in addition to `<NAMESPACE>:group:parent:child:grandchild`.

Ownership of any role always implies membership in that particular (sub)group. However, holding a more specific role in a subgroup does not imply the same role in the parent group. For example:

`<NAMESPACE>:group:parent:child:role=manager`

implies plain membership in both `child` and `parent`, but NOT:

`<NAMESPACE>:group:parent:role=manager`

[5] URN-equivalence: https://tools.ietf.org/html/rfc8141#section-3

2.1 Character encoding

With regard to characters outside the ASCII range, as per [RFC 8141], URNs representing group and role information MUST use only Unicode characters encoded in UTF-8 and further encoded as required by RFC 3986.

Octet 0 (0 hex) should NEVER be used, in either unencoded or percent-encoded form.

The character encoding rules below apply for characters that appear in the URN components following the NAMESPACE component:

- The following characters MUST be percent-encoded using the method defined in Section 2.1 of the generic URI specification [RFC 3986]:
  - Any characters outside the ASCII range
  - The characters ":" and "#" when not used as URN delimiters
  - The question mark character "?" when not used inside the f-component, i.e. the `<AUTHORITY>`
  - The space character (" "). I.e it MUST be encoded as "%20", as opposed to encoding it as a "+".
  - The "=" character when not used within the `role=` literal which introduces the `<ROLE>` component
  - The characters "\", """, "<", ">", "[", "]", "^", "`", "{", "|", and "}"
  - Octets 1-31 (1-20 hex) |
  - Octets 127-255 (7F-FF hex)
- Other characters MUST NOT be percent-encoded

2.2 Normalisation

When constructing URN values expressing group membership and role information case normalisation (as specified in Section 6.2.2.1 of [RFC3986]) MUST be applied to:

1. The <NAMESPACE> component, by conversion to lower case
2. Any hexadecimal digits within a percent-encoding triplet (e.g., "%3a" versus "%3A") in the <GROUP>, <ROLE> and <AUTHORITY> components, by conversion to upper case for the digits A-F.

For example, each of the URN values below:

```
URN:EXAMPLE:FOO:group:Minun%20Ryhm%C3%A4ni
URN:EXAMPLE:foo:group:Minun%20Ryhm%C3%A4ni
URN:example:foo:group:Minun%20Ryhm%C3%A4ni
urn:example:foo:group:Minun%20Ryhm%c3%a4ni
urn:example:foo:group:Minun%20Ryhm%c3%A4ni
```

would be expressed as follows after applying the normalisation process:

```
urn:example:foo:group:Minun%20Ryhm%C3%A4ni
```

2.3 Equivalence

Two URNs are URN-equivalent if their assigned-name portions are octet-by-octet equal[6]. Note that the group authority information specified in the f-component of the URN MUST be ignored in this process. For example, the following URNs are equivalent:

```
<NAMESPACE>:group:parent:role=manager#authority1

<NAMESPACE>:group:parent:role=manager#authority2

<NAMESPACE>:group:parent:role=manager
```

URN components not used by this specification (such as q-components or r-components) SHOULD be ignored when determining equivalence.

[6] Percent-encoding normalisation (as specified in Section 6.2.2.2 of [RFC3986] is not necessary as part of the comparison process, since the character encoding rules in §2.1 already ensure a unique representation.

3 Considerations for different federated identity protocols

This section discusses protocol-specific considerations for communicating group membership and role information.

3.1 Security Assertion Markup Language 2.0 (SAML)

When using SAML, different standardised possibilities are available to convey group membership information. Specifically, both the `isMemberOf` [eduMember] and the `eduPersonEntitlement` [eduPerson] attribute can be used for representing group membership. However, `eduPersonEntitlement` values (formatted as URIs, either URNs or URLs) are in addition used to indicate rights to resources.

Therefore, group membership and role information MUST be communicated to SAML 2.0 relying parties using the `eduPersonEntitlement` attribute following the syntax specified in Section 2 of this document.

It should be noted that while the eduPersonEntitlement is not part of the REFEDS “Research and Scholarship” (R&S) attribute bundle, a relying party may request it if necessary, without violating compliance with the R&S entity category (see [REFEDS-R&S-FAQ]).

3.2 OpenID Connect (OIDC) and OAuth 2.0

In OpenID Connect [OIDC-Core-1] there is no standard claim to carry group membership information. However, SCIM [RFC7643] and [RFC9068] define the multivalued `entitlements` claim for indicating rights to things such as objects or services. The `entitlements` claim specification does not impose any requirements on the vocabulary or syntax.

Therefore group membership and role information MUST be communicated to OpenID Connect/OAuth2 relying parties using at least the `entitlements` claim as defined in [RFC7643], following the syntax defined in Section 2 of this specification. To retain compatibility with [AARC-G002], implementations MAY additionally send the same information using the `eduperson_entitlement` claim.

A relying party can employ any of the following mechanisms to indicate support for the claims defined in AARC-G069 and AARC-G002:

1. Request either the `entitlements` or the `eduperson_entitlement` claim using Scope values as defined in Section 5.4 of the OpenID Connect specification [OIDC-Core-1]. The following[7] Scope values can be used:
    a. `entitlements` - This scope value requests access to the `entitlements` Claim (AARC-G069)
    b. `eduperson_entitlement` - This scope value requests access to the `eduperson_entitlement` Claim (AARC-G002)
2. Request either the `entitlements` or the `eduperson_entitlement` claim using the "claims" Request Parameter as defined in Section 5.5 of the OpenID Connect specification [OIDC-Core-1].
3. Request either the `entitlements` or the `eduperson_entitlement` claim using an out-of-band mechanism (e.g. during relying party registration)

[7] Additional Scope values for requesting these claims may be defined in other AARC Guidelines

References

[AARC-G002] Expressing group membership and role information (AARC-G002); <https://aarc-community.org/guidelines/aarc-g002>

[eduMember] Hazelton, K., "LDAP representations of membership in groups", DOI 10.26869/TI.111.1, July 2005; <http://doi.org/10.26869/TI.111.1>

[eduPerson] Internet2, "eduPerson Object Class Specification (202111)", November 2021; <https://wiki.refeds.org/display/STAN/eduPerson+2021-11>

[OIDC-Core-1] OpenID Connect Core 1.0, <https://openid.net/specs/openid-connect-core-1_0.html>

[REFEDS-R&S-FAQ] REFEDS, "Research & Scholarship Entity Category FAQ"; <https://wiki.refeds.org/display/ENT/Research+and+Scholarship+FAQ>

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

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <https://www.rfc-editor.org/info/rfc3986>

[RFC7643] Hunt, P., Ed., Grizzle, K., Wahlstroem, E., and C. Mortimore, "System for Cross-domain Identity Management: Core Schema", RFC 7643, DOI 10.17487/RFC7643, September 2015, <https://www.rfc-editor.org/info/rfc7643>

[RFC8141] Saint-Andre, P. and J. Klensin, "Uniform Resource Names (URNs)", RFC 8141, DOI 10.17487/RFC8141, April 2017, <https://www.rfc-editor.org/info/rfc8141>

[RFC9068] Bertocci, V., "JSON Web Token (JWT) Profile for OAuth 2.0 Access Tokens", RFC 9068, DOI 10.17487/RFC9068, October 2021, <https://www.rfc-editor.org/info/rfc9068>

[VOMS-AC] "The VOMS Attribute Certificate Format", <https://www.ogf.org/documents/GFD.182.pdf>

Annex A: Example mappings with existing group representation standards

Standard	Original value	Mapped value
VOMS FQAN [VOMS-AC]	*/vo.example.org*	<NAMESPACE>:group:*vo.example.org*
	/vo.example.org/Role=NULL	<NAMESPACE>:group:*vo.example.org*
	/vo.example.org/Role=*manager*	<NAMESPACE>:group:vo.example.org:role=manager
	/vo.example.org/thegroup/thesubgroup/thesubsubgroup	<NAMESPACE>:group:*vo.example.org:thegroup:thesubgroup:thesubsubgroup*
	/vo.example.org/thegroup/thesubgroup/thesubsubgroup/Role=NULL	<NAMESPACE>:group:*vo.example.org:thegroup:thesubgroup:thesubsubgroup*
	/vo.example.org/thegroup/thesubgroup/thesubsubgroup/Role=*manager*	<NAMESPACE>:group:*vo.example.org:thegroup:thesubgroup:thesubsubgroup*:role=*manager*
SCIM [RFC7643]	{ "id": "*8878ae43-965a-412a-87b5-38c398a76569*", "displayName": "Project on group APIs" }	<NAMESPACE>:group:*8878ae43-965a-412a-87b5-38c398a76569*

Annex B: Example access control rules based on group membership and role information

This annex provides examples of rules for controlling access to resources based on the group membership and role information expressed through the URN-formatted values specified in Section 2. The examples below assume a relying party software that supports attribute-based access control using regular expressions. Example configurations for specific relying party software implementations can be found via [AARC-G069].

Example 1: Relying party permitting access to all members of a specific group, named *parentgroup*, under the *urn:example:foo* namespace

```
^urn:example:foo:group:parentgroup(:[^:#]+)*(:role=[^:#]+)?(#.+)?$
```

which would match for example the following URN values:

```
urn:example:foo:group:parentgroup
urn:example:foo:group:parentgroup:role=manager
urn:example:foo:group:parentgroup#authority
urn:example:foo:group:parentgroup:role=manager#authority
urn:example:foo:group:parentgroup:childgroup:role=manager
urn:example:foo:group:parentgroup:childgroup:grandchildgroup:role=manager
```

Example 2: Relying party permitting access to all group members who are assigned the *myrole* role under the *urn:example:foo* namespace:

```
^urn:example:foo:group(:[^:]+)+:role=myrole(#.+)?$
```

which would match for example the following URN values:

```
urn:example:foo:group:mygroup:role=myrole
urn:example:foo:group:mygroup:mysubgroup:role=myrole
urn:example:foo:group:mygroup:role=myrole#authority
```

Example 3: Relying party permitting access to all members of the *mygroup* subgroup under the *urn:example:foo* namespace:

```
^urn:example:foo:group(:[^:]+)*:mygroup(:[^:]+)*(:role=[^:#]+)?(#.+)?$
```

which would match for example the following URN values:

```
urn:example:foo:group:mygroup
urn:example:foo:group:mygroup:role=somerole
urn:example:foo:group:mygroup:anothergroup
urn:example:foo:group:anothergroup:mygroup:role=somerole
urn:example:foo:group:anothergroup:mygroup:role=somerole#authority
```