Guidelines on expressing group membership and role information (201710)

Published Date: 13-10-2017

Version: 201710

Document Code: AARC-JRA1.4A

Comments to: aarc-connect@lists.geant.org

Status of this document

This version of the guidelines for expressing group membership and role information, supersedes version 1.0, which was released on 13-06-2017.

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 730941 (AARC2).

Contents

1 Introduction

Information about the groups a user is a member of is commonly used by SPs 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 managing services. Such services often act as Attribute Authorities, maintaining additional information about the users, including VO membership, group membership within VOs, as well as user roles. It is therefore necessary that all involved SPs 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, both syntactically and semantically:
  - Syntactic: Uniform formatting; for example, representing group membership as URNs within a specific namespace and a set of rules for the NSS portion
  - Semantic: Common representation of equivalent concepts; for instance, "admin" and "manager" should be communicated to end SPs as "manager"
- 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

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

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 RFC 2119 [RFC2119].

2 General guidelines

The guidelines presented in this section, have been defined based on experiences from multiple parties in the AARC project and have subsequently been discussed and tested through the Service Activity 1 Pilots (SA1) attribute management pilot [AARC-SA1-AMP]. Furthermore, it should be noted that a group membership representation scheme following these recommendations has already been adopted to enable cross-infrastructure exchange of group information between the EGI and the ELIXIR AAI.

- **Centralised harmonisation of group membership information**

Adopt a proxy-based AAI, to delegate to the proxy component the complexity of dealing with different group membership representations that originate from diverse IdPs/AAs. As a result, the end SPs will not have to handle the harmonisation of group membership information as this will be performed in a centralised fashion by the SP proxy.

- **Compatibility with existing group information models**

  Adopt a group representation scheme that can be easily translated to/from standardised or widely used group data models, such as SCIM, VOOT or VOMS, and POSIX systems, if required.

- **Scoping of group membership information**

  Specify the scopes where the identified group membership information is valid. These scopes should include:

  - The authoritative source for each piece of group membership information.
  - The VO associated with the identified group.
  - The entire chain of group components, from the root parent group to the identified child group (in the case of group hierarchies).

  The rationale behind scoping is to prevent clashes between groups that are managed by different VOs/administrative domains. This eliminates the need for syntactic 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 SPs for quick authorisation decisions.

- **Use the eduPersonEntitlement attribute**

  When using SAML, different standardised possibilities are available to convey group membership information. Specifically, both the isMemberOf [SWITCH-IMO] and the eduPersonEntitlement attribute [I2-EPE] 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. In the case of OpenID Connect there is currently no standard claim to carry group membership information. However, the REFEDS OpenID Connect for Research and Education Working Group [OIDCre] is already investigating the standardisation of new claims for expressing the attributes defined in the eduPerson schema [I2-EP].

  It should be noted that while the eduPersonEntitlement is not part of the REFEDS "Research and Scholarship" (R&S) [REFEDS-RS] attribute bundle, an SP may request it if necessary [REFEDS-RS-1], without violating compliance with the R&S entity category. However, SPs are still encouraged to stick to the R&S bundle wherever possible.

- **Use of valid URIs for representing group membership information**

  As of 2015, MACE [MACE] encourages the use of URLs in preference to URNs [MACE-SR].

  Benefits of using URLs instead of URNs include:

  - Legitimate URL values are globally unique if a suitable (sub)domain is used and a delegation model is in place for defining paths under that root domain. No one else has the legal right to create values under that (sub)domain, so any assignments made under that subdomain will be globally unique.
  - If the URLs resolve to web pages, it is possible to make the assigned values self-documenting by posting a definition of the value at that URL. In practice, however, the relevant domain that is used for resolvable URLs is often the domain of corporate public relations departments and as such is not easily maintainable by technical staff responsible for the AAI.
  - URLs do not require a formal registration for a subtree, as is required for URNs.

Benefits of using URNs instead of URLs include:

- URNs are currently more commonly used for expressing eduPersonEntitlement values by existing IdPs/AAs/federations.
- URNs can easily support scoping following a hierarchical structure when necessary. Using the namespace identifier registry delegation model, URN values can thus be managed in a distributed fashion by different issuing authorities, communities/VOs, group management systems.

The use of URLs and URNs each has its merits. Given the wide adoption of URNs, these guidelines suggest the use of URNs for formatting eduPersonEntitlement attribute values that express group membership and role information. The specification of the URN namespace for these eduPersonEntitlement values is defined in Section 3.

3 Specification

This section describes a URN namespace for expressing group membership and role information as eduPersonEntitlement values that can be uniformly interpreted across infrastructures.

3.1 Syntax

An eduPersonEntitlement 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>]#<GROUP-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 with IANA [URN-IANA], as per RFC8141, ensuring global uniqueness. Implementers can and should use one of the existing registered URN namespaces, such as `urn:geant` [URN-GEANT] and `urn:mace` [URN-MACE];

  - `<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 an eduPersonEntitlement 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;

- `<GROUP-AUTHORITY>` is a non-empty string that indicates the authoritative source for the entitlement value. For example, it can be the FQDN of the group management system that is responsible for the identified group membership information. The `<GROUP-AUTHORITY>` is specified in the f-component [URN-F-COMP] of the URN; thus, it is introduced by the number sign ("#") character and terminated by the end of the URN. Any characters outside the ASCII range that appear in the `<GROUP-AUTHORITY>` MUST be percent-encoded using the method defined in Section 2.1 of the generic URI specification (RFC 3986 [RFC3986]). As described in Section 3.2, the `<GROUP-AUTHORITY>` MUST NOT be taken into account when determining equivalence [URN-EQUIV] of URN-formatted eduPersonEntitlement values expressing group membership and role information.

3.2 Semantics

Each eduPersonEntitlement 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 member of a subgroup, then they are also member of the parent group. For example:

```
<NAMESPACE>:group:parent-group:child-group#<GROUP-AUTHORITY>
```

implies membership in *`parent-group`*, i.e.:

```
<NAMESPACE>:group:parent-group#<GROUP-AUTHORITY>
```

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-group:child-group:role=manager#<GROUP-AUTHORITY>
```

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

```
<NAMESPACE>:group:parent-group:role=manager#<GROUP-AUTHORITY>
```

Determining if two eduPersonEntitlement values refer to the same group membership (and role, if specified) requires testing for URN-equivalence as per Section 3 of RFC 8141. Thus, the mandatory group authority information specified in the f-component of the URN MUST be ignored in this process. For example, the following two URNs are equivalent to the last example above:

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

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

Annex A: Example mappings with existing group representation standards

Standard	Original value	Mapped value
VOMS FQAN [VOMS-FQAN]	/*vo.example.org*	<NAMESPACE>:group:*vo.example.org*#<GROUP-AUTHORITY>
	/*vo.example.org*/Role=NULL	<NAMESPACE>:group:*vo.example.org*#<GROUP-AUTHORITY>
	/*vo.example.org*/Role=*manager*	<NAMESPACE>:group:*vo.example.org*:role=*manager*#<GROUP-AUTHORITY>
	/*vo.example.org*/*thegroup*/*thesubgroup*/*thesubsubgroup*	<NAMESPACE>:group:*vo.example.org*:*thegroup*:*thesubgroup*:*thesubsubgroup*#<GROUP-AUTHORITY>
	/*vo.example.org*/*thegroup*/*thesubgroup*/*thesubsubgroup*/Role=NULL	<NAMESPACE>:group:*vo.example.org*:*thegroup*:*thesubgroup*:*thesubsubgroup*#<GROUP-AUTHORITY>
	/*vo.example.org*/*thegroup*/*thesubgroup*/*thesubsubgroup*/Role=manager	<NAMESPACE>:group:*vo.example.org*:*thegroup*:*thesubgroup*:*thesubsubgroup*:role=*manager*#<GROUP-AUTHORITY>

SCIM [SCIM] / VOOT [VOOT]	{
	"id": "*8878ae43-965a-412a-87b5-38c398a76569*",
	"displayName": "*Project on group APIs*"
	}
		<NAMESPACE>:group:*8878ae43-965a-412a-87b5-38c398a76569*#<GROUP-AUTHORITY>

VOOT	{
	"id": "*e01eafb1-5f1c-4992-fcd5-ab0160c7ad24*",
	"displayName": "*Course M.201 Mathematics at University of Oslo*",
	"membership": {
	"basic": "*member*",
	}
	}
		<NAMESPACE>:group:*e01eafb1-5f1c-4992-fcd5-ab0160c7ad24*:role=*member*#<GROUP-AUTHORITY>

Standard	Original value	Mapped value

	{
	"id": "*e01eafb1-5f1c-4992-fcd5-ab0160c7ad24*",
	"displayName": "*Course M.201 Mathematics at University of Oslo*",
	"membership": {
	"basic": "admin",
	}
	}
		<NAMESPACE>:group:*e01eafb1-5f1c-4992-fcd5-ab0160c7ad24*:role=*admin*#<GROUP-AUTHORITY>

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 proposed URN-formatted eduPersonEntitlement attribute values. The examples below assume a Shibboleth-based SP but they can easily be adapted to other implementations that support attribute-based access control using regular expressions.

Example 1: SP permitting access to members of a specific group, i.e. *`group1`*, under the *`urn:example:foo`* namespace; the SP only accepts membership information originating from group authority *`group-auth1`*. Using Apache configuration, this can be done as follows:

```
Require shib-attr entitlement urn:example:foo:group:group1#group-auth1
```

Alternatively, using the XML-based Access Control plugin for Shibboleth:

```
<AccessControl type="edu.internet2.middleware.shibboleth.sp.provider.XMLAccessControl">
  <Rule require="entitlement">
    urn:example:foo:group:group1#group-auth1
  </Rule>
</AccessControl>
```

Example 2: SP permitting access to all group members, regardless of role, under two specific namespaces, *`urn:example:foo`* and *`urn:example:bar`*:

```
<AccessControl type="edu.internet2.middleware.shibboleth.sp.provider.XMLAccessControl">
  <OR>
    <RuleRegex require="entitlement">^urn:example:foo:group:.*$</RuleRegex>
    <RuleRegex require="entitlement">^urn:example:bar:group:.*$</RuleRegex>
  </OR>
</AccessControl>
```

Example 3: SP permitting access to all group members under a given namespace, matching two specific attribute authorities, *`group-auth1`* and *`group-auth2`*:

```
<AccessControl type="edu.internet2.middleware.shibboleth.sp.provider.XMLAccessControl">
  <OR>
    <RuleRegex require="entitlement">
      ^urn:example:foo:group:.*#group-auth1$
    </RuleRegex>
    <RuleRegex require="entitlement">
      ^urn:example:foo:group:.*#group-auth2$
    </RuleRegex>
  /OR>
</AccessControl>
```

Example 4: SP permitting access to all group members from specific group authorities:

```
<AccessControl type="edu.internet2.middleware.shibboleth.sp.provider.XMLAccessControl">
  <OR>
    <RuleRegex require="entitlement">^urn:.*#group-auth1$</RuleRegex>
    <RuleRegex require="entitlement">^urn:.*#group-auth3$</RuleRegex>
  /OR>
</AccessControl>
```

Example 5: SP permitting access to all group members who are assigned the `manager` role under specific namespaces, regardless of the group authority:

```
<AccessControl type="edu.internet2.middleware.shibboleth.sp.provider.XMLAccessControl">
  <OR>
    <RuleRegex require="entitlement">
      ^urn:example:foo:group:.*:role=manager#.*$
    </RuleRegex>
    <RuleRegex require="entitlement">
      ^urn:example:bar:group:.*:role=manager#.*$
    </RuleRegex>
  /OR>
</AccessControl>
```

References

[AARC-SA1-AMP] Attribute Management Pilot wiki
https://wiki.geant.org/display/AARC/AttributeManagementPilot

[I2-EP] eduPerson Object Class Specification (201310)
http://software.internet2.edu/eduperson/internet2-mace-dir-eduperson-201310.html

[I2-EPE] eduPersonEntitlement description
http://software.internet2.edu/eduperson/internet2-mace-dir-eduperson-201310.html#eduPersonEntitlement

[MACE] MACE website
https://www.internet2.edu/communities-groups/middleware/middleware-architecture-committee-education-mace/

[MACE-SR] "Information for organisations requesting a delegated namespace"
https://www.internet2.edu/products-services/trust-identity/mace-registries/#service-registries

[OIDCre] REFEDS OpenID Connect for Research and Education Working Group
https://wiki.refeds.org/display/GROUPS/OIDCre

[REFEDS-RS] REFEDS web page: Research and Scholarship Entity Category
https://refeds.org/category/research-and-scholarship

[REFEDS-RS-1] REFEDS wiki page: "Are SPs allowed to request attributes other than R&S attributes?"
https://wiki.refeds.org/display/ENT/Research+and+Scholarship+FAQ#ResearchandScholarshipFAQ-AreSPsallowedtorequestattributesotherthanR&Sattributes?

[RFC2119] Key words for use in RFCs to Indicate Requirement Levels
https://www.ietf.org/rfc/rfc2119.txt

[RFC3986] Uniform Resource Identifier (URI): Generic Syntax:
https://www.ietf.org/rfc/rfc3986.txt

[SCIM] System for Cross-domain Identity Management
https://tools.ietf.org/html/rfc7644

[VOOT] VOOT
http://openvoot.org/datamodel/

[SWITCH-IMO] isMemberOf description
https://www.switch.ch/aai/support/documents/attributes/ismemberof/index.html

[URN-F-COMP] URN f-component:
https://tools.ietf.org/html/rfc8141 - section-2.3.3

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

[URN-GEANT] GÉANT URN Namespace:
https://www.geant.org/Services/Trust_identity_and_security/Pages/NamespaceRegistry.aspx

[URN-IANA] IANA Registry of URN Namespaces:
https://www.iana.org/assignments/urn-namespaces/urn-namespaces.xhtml

[URN-MACE] MACE URN Namespace
https://www.internet2.edu/products-services/trust-identity/mace-registries/urnmace-namespace/

[VOMS-FQAN] VOMS Fully Qualified Attribute Name:
https://www.ogf.org/documents/GFD.182.pdf

Glossary

AA	Attribute Authority
AAI	Authentication and Authorisation Infrastructure
EGI	European Grid Infrastructure
FQDN	Fully Qualified Domain Name
IdP	Identity Provider
MACE	Middleware Architecture Committee for Education
NSS	Namespace Specific String
OIDCre	OpenID Connect for Research and Education
POSIX	Portable Operating System Interface
REFEDS	Research and Education FEDerations group
R&S	Research and Scolarship
SCIM	System for Cross-domain Identity Management
SP	Service Provider
URL	Uniform Resource Locator
URI	Uniform Resource Identifier
URN	Uniform Resource Name
VO	Virtual Organization
VOMS	Virtual Organization Membership Service
VOOT	Virtual Organization Orthogonal Technology