CLARIN Federated Content Search (CLARIN-FCS) - AAI 1.0

Oliver Schonefeld, Leif-Jöran Olsson, Thomas Eckart, André Moreira, Erik Körner

Version 1.0, 2023-05-11

Table of Contents

Chapter 1. Introduction

This specification is an extension specification to the CLARIN-FCS Core 2.0 specification and describes an access mechanism for restricted resources.

The Federated Content Search currently does not support the restriction of access to resources to specific user groups or users. There is no mechanism in the aggregator that limits access to announced resources and endpoints are typically accessible via the FCS protocol directly (omitting the use of the FCS aggregator). The goal is to allow restricting access to FCS resources to users that are authenticated using the established AAI infrastructure.

This contains the following issues:

1. Shibbolizing the aggregator frontend so that
   - unauthenticated users still get access to all unrestricted endpoints
   - authenticated users get access to all unrestricted endpoints and additionally to the restricted endpoints
2. Specifying (and implementing) changes for endpoints so that:
   - endpoints can announce restricted resources to the aggregator
   - endpoints can rely on authenticated FCS user requests

1.1. Terminology

The key words `MUST`, `MUST NOT`, `REQUIRED`, `SHALL`, `SHALL NOT`, `SHOULD`, `SHOULD NOT`, `RECOMMENDED`, `MAY`, and `OPTIONAL` in this document are to be interpreted as in RFC2119.

1.2. Glossary

AAI

Authentication & Authorization Infrastructure. A service and a procedure that enables members of different institutions to access protected information that is distributed on different web servers. See Shibboleth.

IdP

Identity Provider, a system entity that issues authentication assertions, see SAML and AAI.

JWT

JSON Web Token, see RFC7519.

RSA

Asymmetric public-key cryptography system by *Rivest-Shamir-Adleman* for digital signatures and encryption.

SAML

Security Assertion Markup Language, an open standard for exchanging authentication and

authorization data.

Shibboleth

Shibboleth is a single sign-on log-in system for computer networks and the Internet. See also SAML.

SP

Service Provider, a system entity that receives and accepts authentication assertions, see SAML and AAI.

1.3. Normative References

RFC2119

Key words for use in RFCs to Indicate Requirement Levels, IETF RFC 2119, March 1997, https://www.ietf.org/rfc/rfc2119.html

RFC7519

JSON Web Token (JWT), IETF RFC 7519, May 2015, https://www.ietf.org/rfc/rfc7519.html

XML-Namespaces

Namespaces in XML 1.0 (Third Edition), W3C, 8 December 2009, http://www.w3.org/TR/2009/REC-xml-names-20091208/

LOC-DIAG

SRU Version 1.2: SRU Diagnostics List, Library of Congress, http://www.loc.gov/standards/sru/diagnostics/diagnosticsList.html

CLARIN-FCS-Core 2.0

CLARIN Federated Content Search (CLARIN-FCS) - Core 2.0, SCCTC FCS Task-Force, May 2017, https://trac.clarin.eu/wiki/FCS/Specification, https://office.clarin.eu/v/CE-2017-1046-FCS-Specification-v89.pdf

1.4. Non-Normative References

RFC6838

Media Type Specifications and Registration Procedures, IETF RFC 6838, January 2013, https://www.ietf.org/rfc/rfc6838.txt

RFC3023

XML Media Types, IETF RFC 3023, January 2001, https://www.ietf.org/rfc/rfc3023.txt

RFC8017

PKCS #1: RSA Cryptography Specifications Version 2.2, IETF RFC 8017, November 2016, https://www.ietf.org/rfc/rfc8017.txt

1.5. Typographic and XML Namespace conventions

The following typographic conventions for XML fragments will be used throughout this specification:

- `<prefix:Element>`

  An XML element with the Generic Identifier *Element* that is bound to an XML namespace denoted by the prefix *prefix*.

- `@attr`

  An XML attribute with the name *attr*.

- `string`

  The literal *string* must be used either as element content or attribute value.

Endpoints and Clients `MUST` adhere to the XML-Namespaces specification. The CLARIN-FCS interface specification generally does not dictate whether XML elements should be serialized in their prefixed or non-prefixed syntax, but Endpoints `MUST` ensure that the correct XML namespace is used for elements and that XML namespaces are declared correctly. Clients `MUST` be agnostic regarding syntax for serializing the XML elements, i.e. if the prefixed or un-prefixed variant was used, and `SHOULD` operate solely on expanded names, i.e. pairs of namespace name and local name.

Chapter 2. Restricted Access to Resources

The FCS supports restriction of access to resources to specific user groups or users that are authenticated using the established AAI infrastructure while still allowing unauthenticated access to all unrestricted resources. This mechanism limits access to announced resources that are typically accessible via the FCS protocol directly (omitting the use of the FCS aggregator).

2.1. Technical Description

The aggregator allows an optional login via Shibboleth. Authentication of search queries which are sent to an FCS endpoint is implemented using authentication headers with JSON Web Token (JWT). User information (e.g. mail address or similar) is encoded in the token as claims.

JWTs include an RSA signed token which `SHOULD` be checked by the respective endpoint. In case of missing or insufficient authorization when accessing restricted resources, the endpoint `MUST` rise an error using SRU diagnostic "Authentication error" (`info:srw/diagnostic/1/3`).

The aggregator owns a single common RSA key for all endpoints. For all endpoints with restricted resources its public key has to be manually transferred and included in their configuration. Communication between aggregator and endpoints is only allowed via SSL. Claims *iss*, *sub* and *aud* `SHOULD` be encoded in the JWT. The endpoint `MUST` check *aud* to see if it is the correct recipient and `MAY` check the fields *iss* and *exp*.

The availability and nature of a personal identifier attribute in the aggregator itself, is not guaranteed for a successful login. In SAML this is limited by the configuration of the external IdP selected by the user. When available, one of three SAML attributes is used as user personal identifier: *email*, *eduPersonPrincipalName* or *eduPersonTargetedID*. These attributes are mapped by the aggregator to *userID* in the same order of preference. *userID* is then passed to the endpoint. The aggregator issues this information to JWTs and sends it to the appropriate endpoints. The endpoints alone announces which authentication information is required (see next section) and decides at runtime whether the information provided is sufficient for access. While the aggregator front end can inform and guide the user in advance, in case authentication information is still missing.

2.2. Announcing Restricted Resources by the Endpoint

In order to announce restricted resources the endpoint needs to explain in the `<Resources>` section of the `<EndpointDescription>` that it does support the aforementioned procedure for a resource. This is done via a dedicated element `<AvailabilityRestriction>`:

Example of AvailabilityRestriction in Endpoint Description for a Resource

```
<Resources>
  <Resource>
    <!-- ... -->
    <Languages>
      <Language>swe</Language>
      <Language>deu</Language>
    </Languages>
```

```
    <AvailabilityRestriction>
      <RestrictionRequirement>requirementName</RestrictionRequirement>
    <!-- ... -->
    </AvailabilityRestriction>
    <!-- ... -->
  </Resource>
</Resources>
```

The `<AvailabilityRestriction>` is defined in the top-level `<Resource>` element and the restriction is valid for its complete sub-tree hierarchy of `<Resource>` elements. From the backwards compatibility perspective this means that if you do not define the `<AvailabilityRestriction>` element all resources will be available for searching through the endpoint.

The text node `requirementName` in `<RestrictionRequirement>` is either of `authOnly` or `personalIdentifier`. `authOnly` does not require any attributes except for the authentication via home institution. If `personalIdentifier` is stated *eduPersonPrincipalName* or *eduPersonTargetedID* or *email* is passed to the endpoint as `userID` using the procedure described above.

Changelog

2023-06-12 — Conversion to AsciiDoc and Migration of specification documents to Github

- Convert specification documents for FCS Core 2.0, Core 1.0, DataView and AAI to AsciiDoc
- Migrate from CLARIN Trac to CLARIN Github
- Add Github Actions workflow to automate build process

2021-01-27 — FCS Update Proposal for AAI

2020-02-19 — Meeting about Shibbolizing FCS