Building an interoperable, distributed storage and authorization system

S. Bertocco[1], B. Major[2], P.Dowler[2], S.Gaudet[2], M.Molinaro[1], G.Taffoni[1]

[1]INAF-OATs: Istituto Nazionale di Astrofisica – Osservatorio Astronomico di Trieste
[2]CADC, National Research Council Canada

Abstract

A joint project between the Canadian Advanced Network for Astronomical Research (CANFAR) and the INAF-Osservatorio Astronomico di Trieste (INAF-OATs), partially funded by the EGI-Engage H2020 European Project, is working to deploy an integrated infrastructure, based on International Virtual Observatory Alliance (IVOA) standards, to access and exploit astronomical data.
CANFAR provides scientists with an access, storage and computation facility, based on software libraries implementing a set of standards developed by the IVOA.
The deployment of a twin infrastructure, basically built on the same open source software libraries, available at https://github.com/opencadc, has been started at INAF-OATs.
At present, this infrastructure provides users with an Access Control Service and a Storage Service based on the VOSpace 2.1 IVOA standard.
The implementation of the IVOA standard ensures the interoperability of the whole geographically distributed storage service.
The Access Control Service is based on the Group Management Service developed at CANFAR and open source available. In the scope of the collaboration, the CANFAR software has been modified to allow integrated user authentication, i.e. users of one of the infrastructures can use resources located either at CANFAR or at INAF-OATs with a single-sign-on access point.
The Storage Service is based on the open source implementation, provided by CANFAR, of the IVOA 2.1 VOSpace standard plus an open source data transfer management service developed at OATs-INAF in the scope of the collaboration.
This poster focuses on the technical choices and the implemented solutions.

IVOA standards

The Virtual Observatory (VO) is the vision that astronomical datasets and other resources should work as a seamless whole, exploitable in a single transparent system.

The International Virtual Observatory Alliance (IVOA) is an organisation that debates and agrees the technical standards that are needed to make the VO possible. It is a framework for discussing and sharing VO ideas and technology, and body for promoting and publicising the VO.

The infrastructure we are describing is based on IVOA standards/recommendations.
In particular:

- CDP 1.0: the Credential Delegation Protocol allows a client program to delegate a user's credentials to a service such that that service may make requests of other services in the name of that user.
- UWS 1.1: the Universal Worker Service Pattern defines how to manage asynchronous execution of jobs on a service.
- VOSI 1.1: IVOA Support Interfaces describes the minimum interface that a web service requires to participate in the IVOA, i.e. a set of common basic functions that all these services should provide in the form of a standard support interface in order to support the effective management of the VO
- VOSpace 2.1: VOSpace is the IVOA interface to distributed storage.
- SSO 1.0.1: IVOA Single-Sign-On Profile describes approved client-server authentication mechanisms

The basic implementation of this standards/recommendations can be found on **https://github.com/opencadc**

An overall view

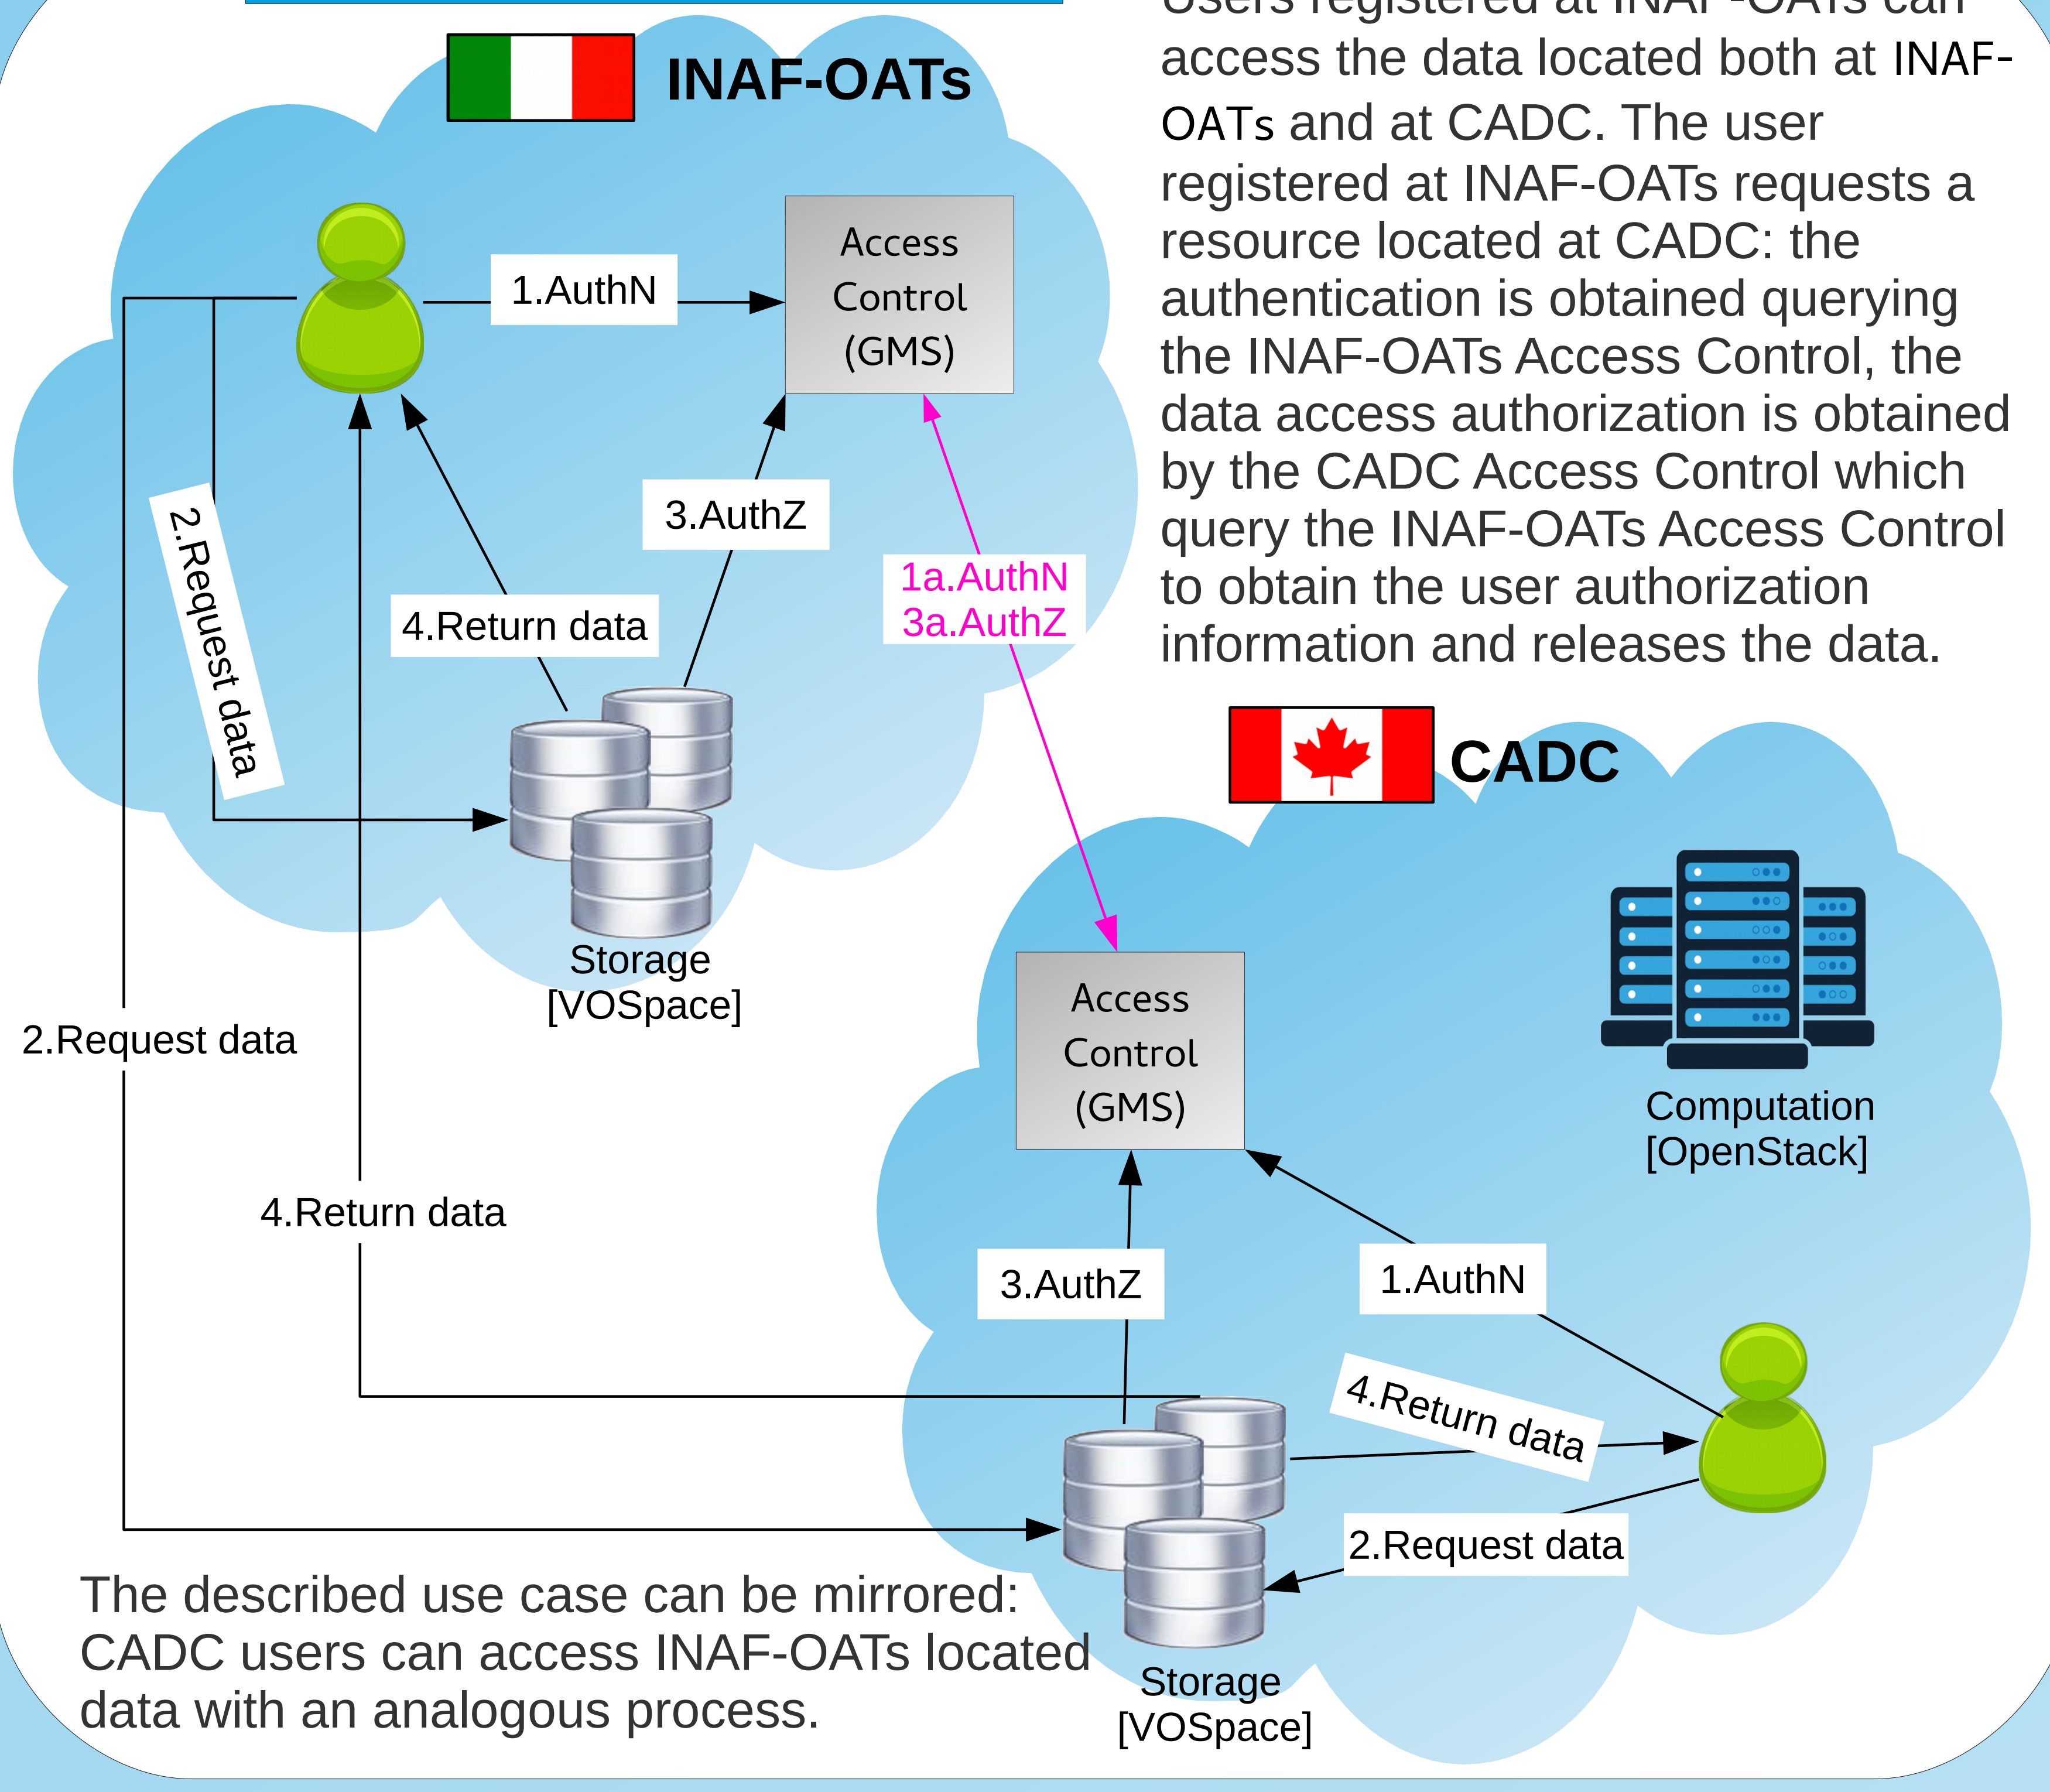

The described use case can be mirrored: CADC users can access INAF-OATs located data with an analogous process.

Users registered at INAF-OATs can access the data located both at INAF-OATs and at CADC. The user registered at INAF-OATs requests a resource located at CADC: the authentication is obtained querying the INAF-OATs Access Control, the data access authorization is obtained by the CADC Access Control which query the INAF-OATs Access Control to obtain the user authorization information and releases the data.

Access Control

Access control is provided through a RESTful service that allows for the management of groups and users, user authentication, and authorization checks by group membership.

The Access Control service provides authentication and identity management support to other web services or to clients directly. Users may have multiple identities and can connect to the ecosystem of services with any of those identities.

Users can be members, administrators, or owners of groups. Ownership and administrative membership allows for different levels of group management.

Users are considered to be authorized to a resource (a service or proprietary data, for example) if they are a member of the group(s) protecting that resource. Resource protection is achieved by the owner of the resource assigning (granting) a group to that resource.

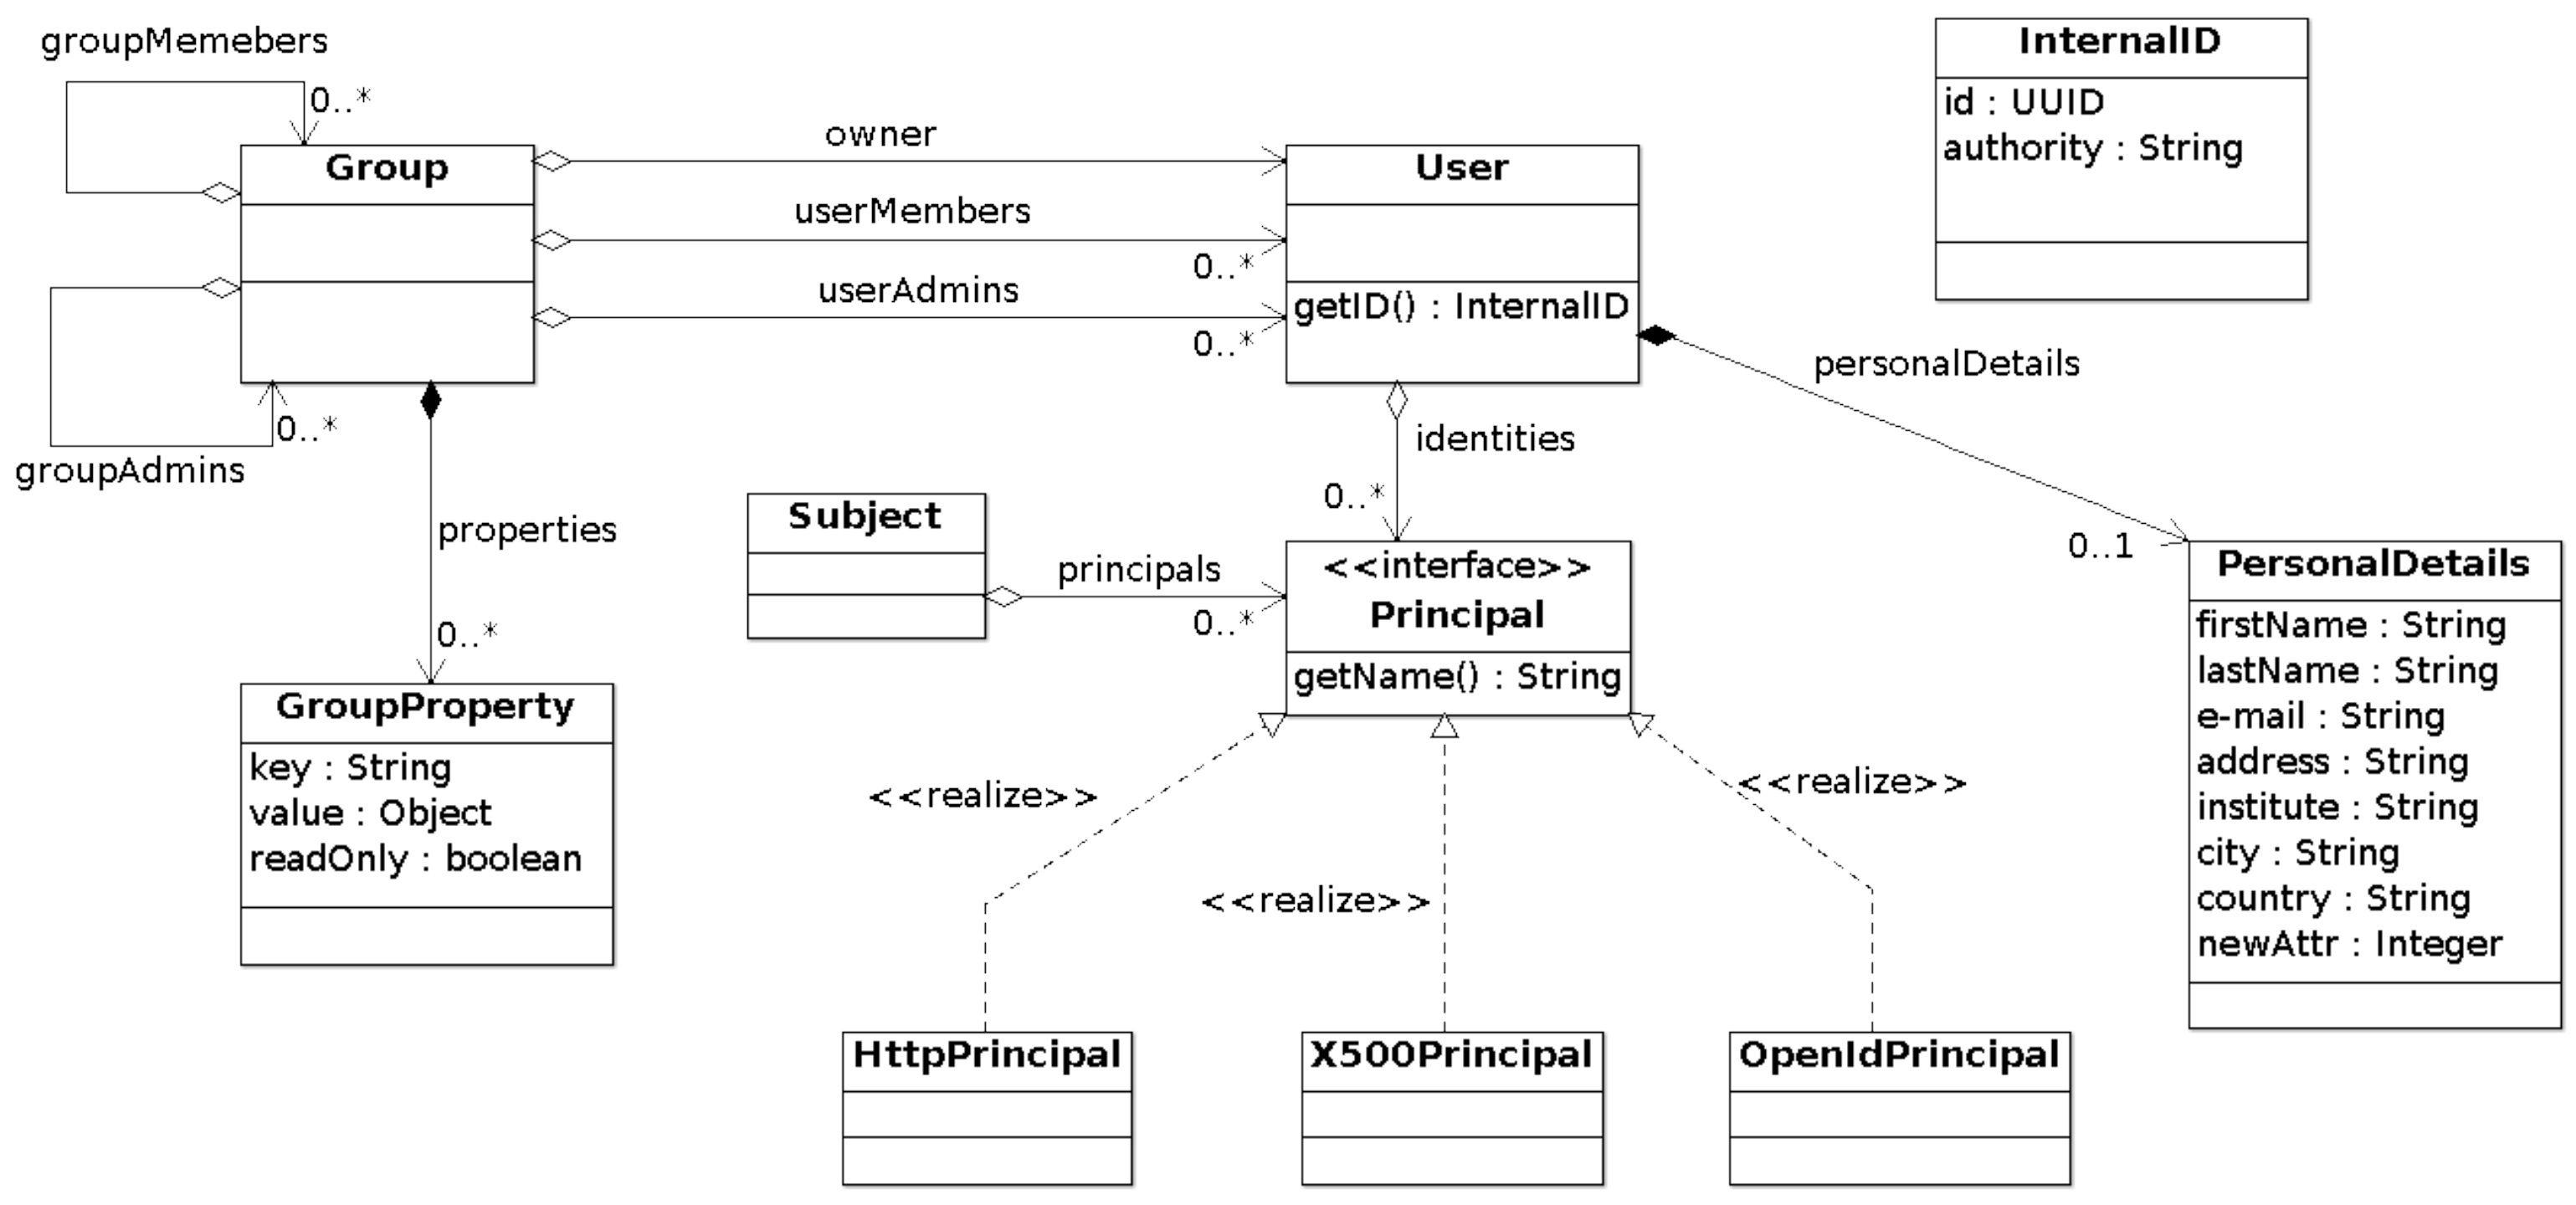

VOSpace Service

The VOSpace is built of two main components: a VOSpace front-end, managing the user requests, the user authentication and authorization and the metadata about the stored data; a VOSpace back-end, managing the physical data storage and retrieval.
The separation of the two components is preserved to allow the use of different back-ends with a minimal effort. Two main concepts are used:

- RestFul architecture (=> easy services distribution)
- Plugins architecture (=> easy technology substitution)

Front-end and back-end expose both a RestFul interface. The communication between the two sides is performed using a TransferGenerator pluggable and configurable component. The Authentication can be customized also implementing an Authorizer interface.

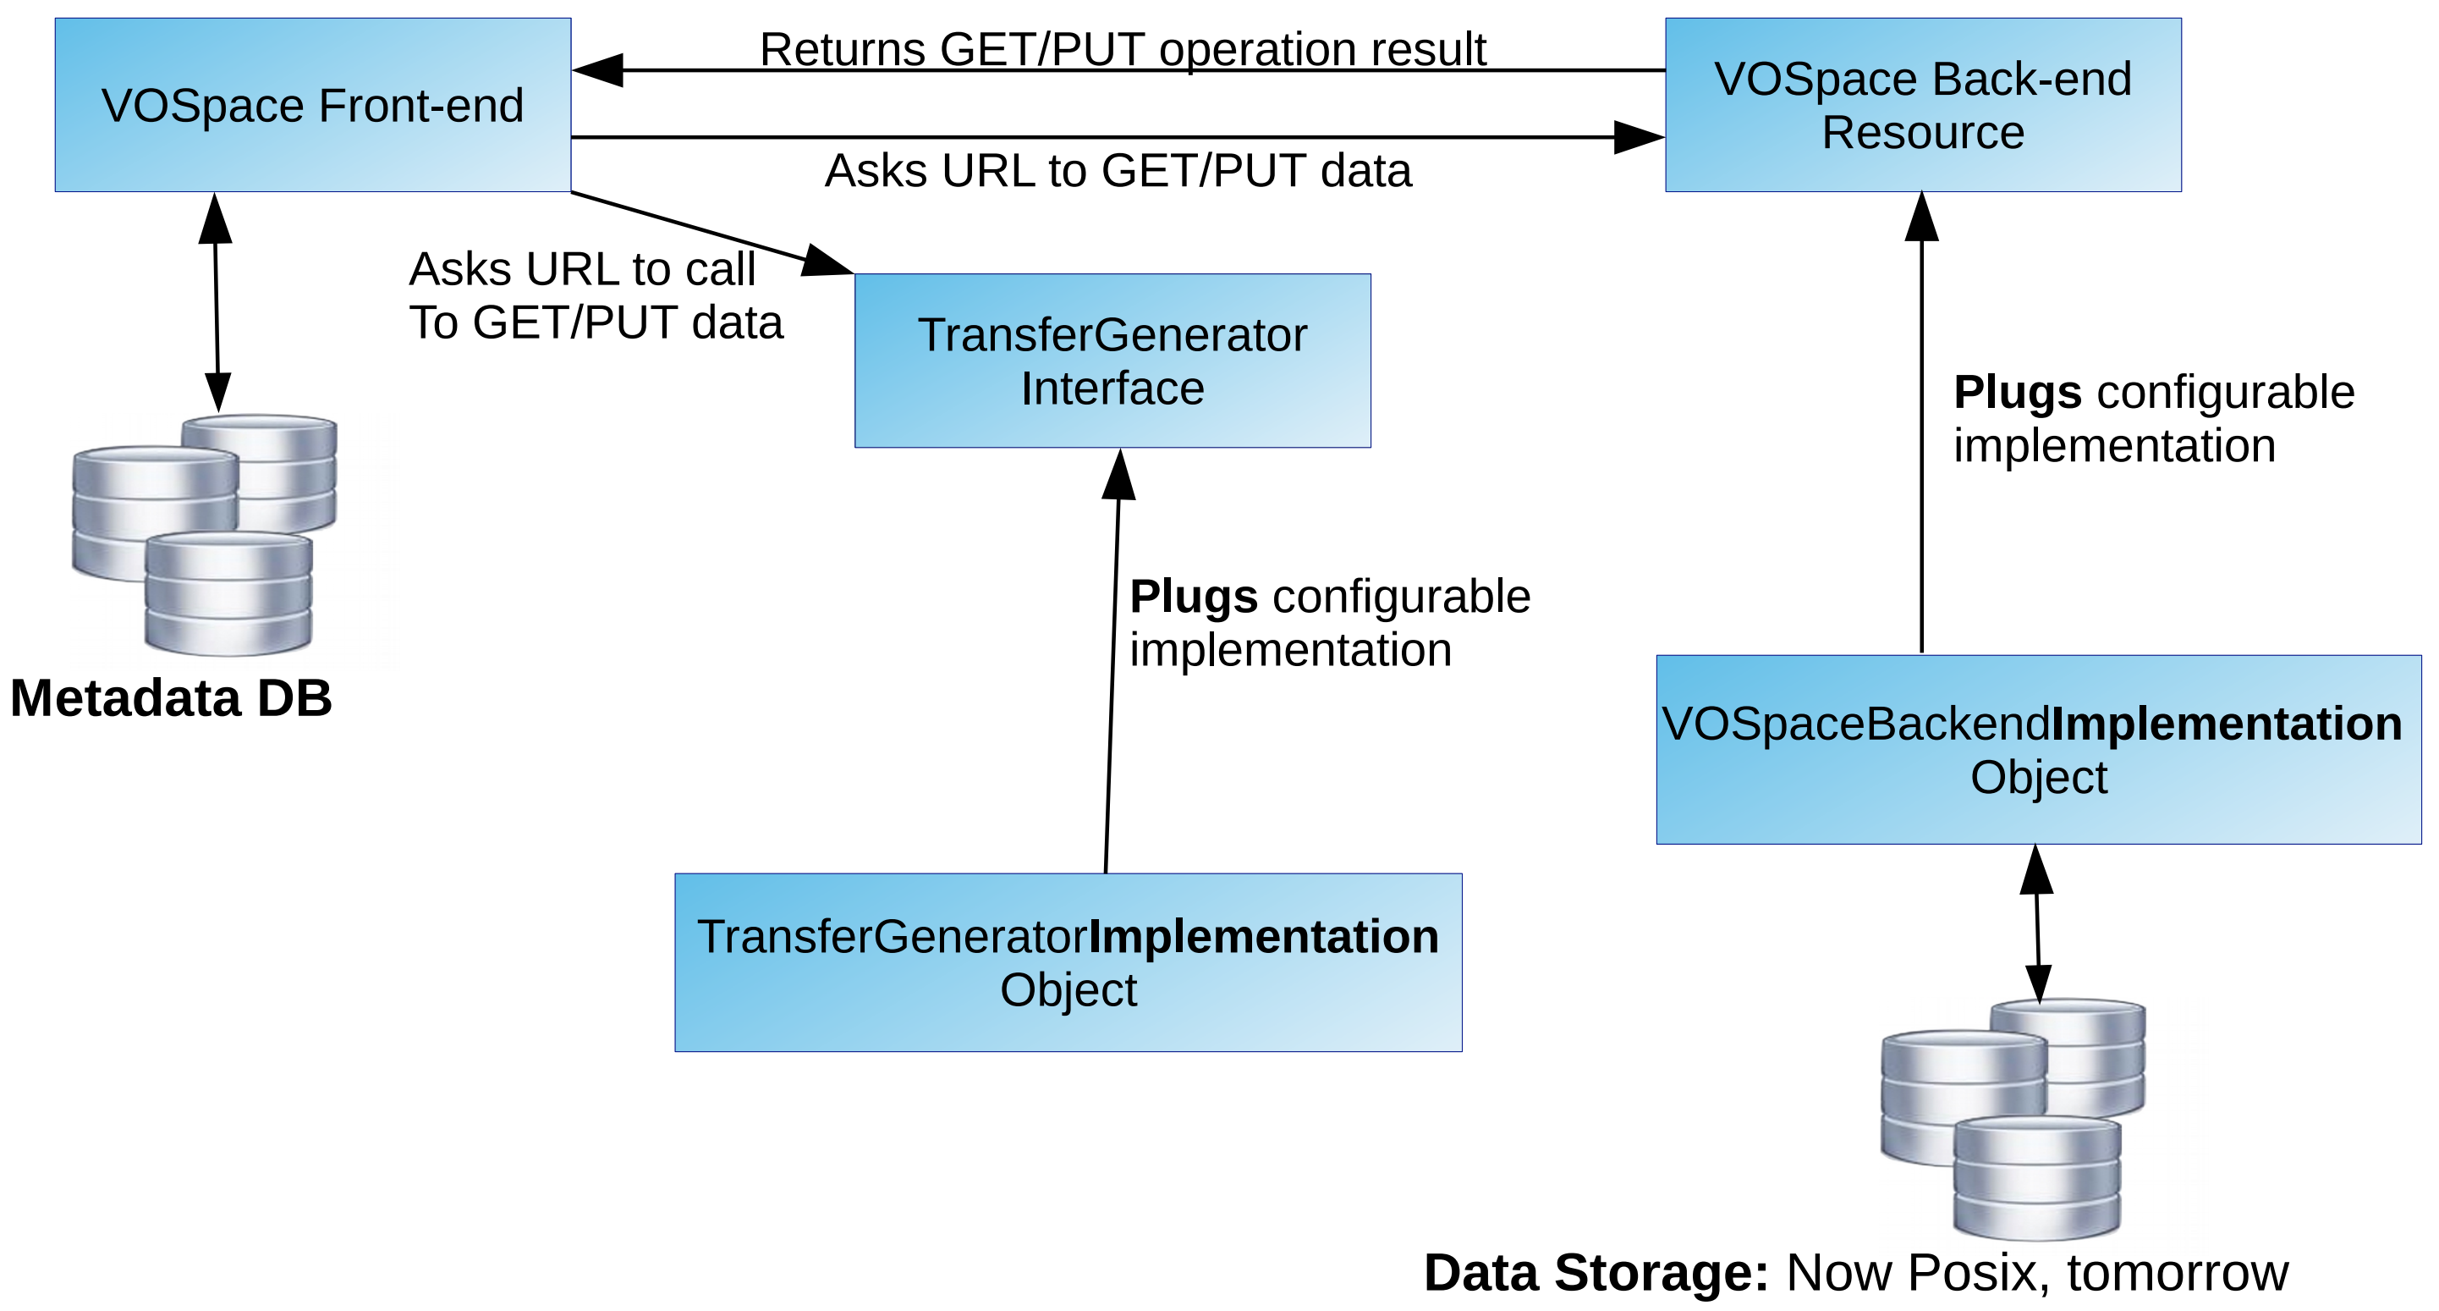
VOSpace Front-end
Returns GET/PUT operation result
Asks URL to GET/PUT data
VOSpace Back-end Resource
Asks URL to call To GET/PUT data
TransferGenerator Interface
Metadata DB
Plugs configurable implementation
Plugs configurable implementation
VOSpaceBackendImplementation Object
TransferGeneratorImplementation Object

Data Storage: Now Posix, tomorrow OpenStack-Swift or Onedata or ...

References and contacts

References:

- www.ivoa.net
- http://wiki.ivoa.net/internal/IVOA/InteropOct2015GWS/InteroperableAA.pdf
- https://github.com/opencadc
- https://github.com/bertocco/opencadc_doc/tree/master/ACUsersGuide

Contacts:

bertocco@oats.inaf.it, brian.major@nrc-cnrc.gc.ca, patrick.dowler@nrc-cnrc.gc.ca, severin.gaudet@nrc-cnrc.gc.ca, molinaro@oats.inaf.it, taffoni@oats.inaf.it