

A Generic Broker Portal linking Authentication and Authorization Infrastructures and Resources

Marc-Alain Steinemann [steine@iam.unibe.ch]

Torsten Braun [braun@iam.unibe.ch]

Institut für Informatik und Angewandte Mathematik, Universität Bern, Neubrückstr.
10, CH-3012 Bern

<http://www.iam.unibe.ch/~rvs>

Abstract. Generating and maintaining user access to resources like e-learning courses is a time and money consuming process for both, users and resource providers. In authentication and authorization infrastructures such as Shibboleth, PAPI and Liberty Alliance, home organizations authenticate their own users and provide a set of user information attributes to the resources they access, depending on the attribute release policies of the user and the home organization. Resources decide by means of the received user information attributes if a user gets access or not. A disadvantage of these infrastructures is that resource interfaces have to be adapted to each of the infrastructures and kept up to date. We propose to fill the gap between authentication and authorization infrastructures and resources with a generic web portal that acts as a broker between authentication and authorization infrastructures and resources. All portal-enabled resources profit from the implemented interfaces to authentication and authorization infrastructures and resources as well as from advanced user and resource management features. The proposed portal has been implemented and connected to the Internet2 middleware called Shibboleth and several types of resources. The software is open source and available for free.

Keywords: Authentication, authorization, resource management, web portal, broker.

1 List of Topics

Introduction

Related Work

Modular AAI Portal Architecture

User and Resource Management

Value Added Services

Discussion and Conclusions

Outlook

Acknowledgments

References

2 Introduction

In recent times the term Authentication and Authorization Infrastructures (AAI) has been used to describe middleware systems consisting of a set of protocols that allow the delegation of authentication and authorization issues to different instances. Authentication is executed by the user's home organization. Authorization is executed in a first step by the home organization and by the respective user by releasing user information attributes towards the querying resource and in a second step by the resource the user wants to access. A resource grants access to a user if the user information attributes meet the expectations of the resource provider. Authentication and authorization infrastructures provide all the necessary mechanisms to enable users, organizations and resources to participate in the system. By these means, authentication and authorization infrastructures connect user communities to resources, for example students to their e-learning courses. Users belong to at least one home organization, for example a university or a college. Resources are systems that provide media content and can be e-learning courses or content management systems. A simplistic overview is given in Figure 1.

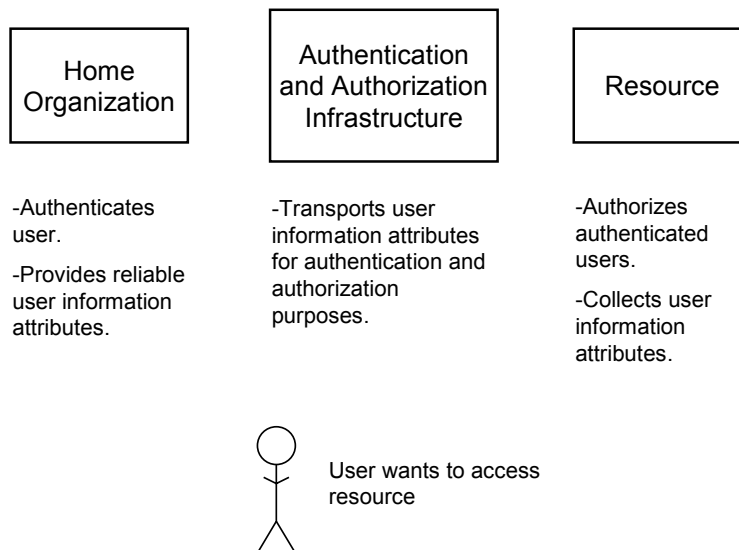


Figure 1: User authentication and information flow.

The biggest advantages resource providers experience by joining an authentication and authorization infrastructures lie in the reduced administrative overhead and the feed of reliable user information in the form of user information attributes, provided by the user's home organization. Resources define their access policy based upon required user

information attributes additionally to user identities. The detailed user management, such as access levels or fine granulated access structures, remains a resource internal issue. Consider the example that user Alice, belonging to a home organization named University of Bern tries to access a resource called TCP/IP course. Alice's home organization as well as the resource is connected to the same AAI. University of Bern stores information attributes about Alice that are released to the resource at the moment Alice tries to access it. Only those attributes are released that are allowed by University of Bern's and Alice's attribute release policy. The TCP/IP course resource receives those attributes and compares them with its own attribute acceptance policy, which describes which user information attributes are required so that the user Alice can access the resource. If Alice's home organization and Alice provide the adequate attributes she can now access the TCP/IP course.

The biggest advantage home organizations experience by joining an authentication and authorization infrastructures lies in the immediate increase of accessible resources by their members. Authentication and authorization infrastructure users get subscribed by their respective home organizations and have to agree with the local subscription policy. Home organizations define their information release policy to protect their members' privacy on the one hand, and on the other hand enable members to access resources. Resource providers decrease administrative overhead as they do not have to manually verify and subscribe users to their own AAI-connected resources.

The biggest advantage users experience by being a member of an authentication and authorization infrastructures lies in a simplified resource access procedure. No on-site or mail registration procedure is needed. It is a precondition in authentication and authorization infrastructures that user information released by home organizations is reliable and accepted by the resource providers. The privacy sphere is guaranteed in two ways: First, users always and only authenticate with their own home organizations and never with resources. Second, users determine, which information attributes about themselves are released to the resource they want to access.

A typical authentication and authorization infrastructure architecture environment as described above with non-AAI-enabled and AAI-enabled resources most likely consists of the elements shown in Figure 2. Additionally to the typical elements also non-AAI-enabled resources and the AAI portal are shown.

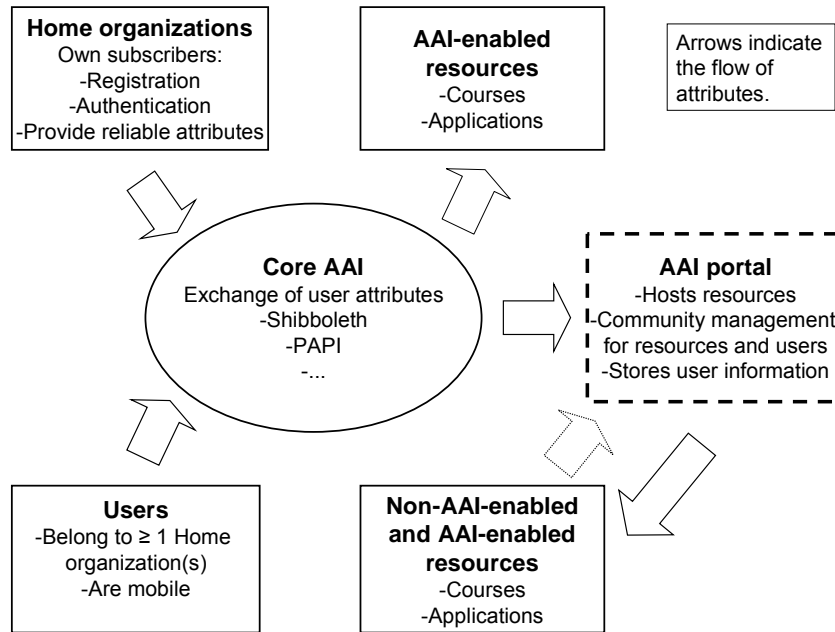


Figure 2: Typical AAI environment with the AAI portal.

A use case for the AAI portal can be found at University of Bern. It grants access to the Virtual Internet and Telecommunications Laboratory of Switzerland (VITELS) [1]. VITELS features innovative e-education for an international audience, such as people from universities, commerce, industry, home sector and less developed countries. VITELS is a complete course covering theory and practice in the area of computer networks with a strong focus on hands-on experience. VITELS students program and configure network devices in exactly the same manner as they would do in a conventional laboratory but now wherever they please. The modular course builds on computational grid architecture and consists of numerous independent, closed learning units comprising theory, practice and tests, contributed and locally operated by the VITELS partners who share time and knowledge. Since 2001, more than 400 students from all over the world have made great use of this highly personalized learning experience.

The grid-like architecture with distributed e-learning laboratories and a central laboratory reservation system caused a high administrative workload before connecting it to the AAI portal. Students sent Emails to info@vitels.ch or tutors sent class lists to us. These data had to be added to the university's students directory and students had to be informed about the access procedure and their credentials. Doing so, VITELS used resources from the IT department and VITELS administrators.

In 2003 we connected VITELS to the AAI portal. We did not natively enable VITELS to the AAI as each AAI update would cause code adaptation on our side but wrote a small adapter for the AAI portal. As the AAI portal and its interface to the AAI is maintained by the national research network and the AAI portal run by the university, we got rid of most administrative work. Most Swiss universities are meanwhile connected to the Swiss-AAI and students get access to the AAI portal where VITELS is hosted. Students need two clicks to subscribe and a VITELS administrator grants or rejects the submission by viewing

the user data. We also profit from the additional possibility to open local user accounts on the AAI portal for users without AAI account.

In section 5 we describe a middleware architecture and an implementation of a solution for user and resource management for all types of resources and authentication and authorization infrastructures. The presented solution helps to overcome open issues especially for non-AAI-enabled resources. In section 8 we present the user and resource management features that allow the collection of additional user information and provide advanced notification mechanisms. The architecture is still extensible in many directions such as accounting, personalized user resources or content adaptation as partly presented in section 6. The portal uses a plug-in concept for resource adaptors, providing a high reusability of any existing resource adaptors. It has been implemented and licensed under the general public license and been transferred to SWITCH [2]. SWITCH, the national research network published the code on SourceForge [3] together with the full documentation and a demonstration implementation.

3 Related Work

Northern American's Internet2 authentication and authorization infrastructure initiative called Shibboleth [4] enables resources to be connected as soon as they have installed and configured Shibboleth's target site software together with the respective security certificates. Home organizations have to install Shibboleth's origin site software and connect their user database to it. For a resource provider this means installing Shibboleth target site and server certificates. Shibboleth does not advertise or list connected resources. Shibboleth is based on federated user management (i.e. user management is delegated to home organizations).

Similar to Shibboleth is the Spanish RedIRIS' [5] authentication and authorization infrastructure initiative called Point of Access to Providers of Information (PAPI) [6]. PAPI tries to keep authentication to the user's home organization and authorization to the resources. After authenticating at their home organization's authentication server, users get a list of all eligible resources. PAPI is based on federated user management.

Liberty Alliance [7] commenced in 2001 and aims to provide open standards for a trust network with a strong emphasis on commercialization. Accounting is a must for commercialization and raises many questions about privacy and security. Liberty Alliance is based on federated user management.

Microsoft .NET Passport [8] has central databases where all sensitive user data is stored. Only one company administrates the stored user data. All resources have to query the same database. Standards are not open and thus make the system insecure compared to open source systems where the code is freely available for anybody to verify. Users cannot define their information release policy to each resource they visit. Several severe security issues have been discovered in the past.

4 Modular AAI Portal Architecture

The AAI portal can be imagined as an entity between authentication and authorization infrastructures and resources as shown in Figure 3. Examples for AAI are Shibboleth, PAPI and Liberty Alliance. At the time of writing this article, only the interface to Shibboleth has been implemented. The interfaces depicted as plugs and connectors have to fit to each other. The advantage of this plug-in concept lies in the reusability of the interfaces. Once an

interface (known as resource adaptor) has been implemented, other similar resources can reuse a resource adaptor or adapt it accordingly to their own needs. On the right side four example resources are shown. In this case, each of the resources uses a different resource adaptor.

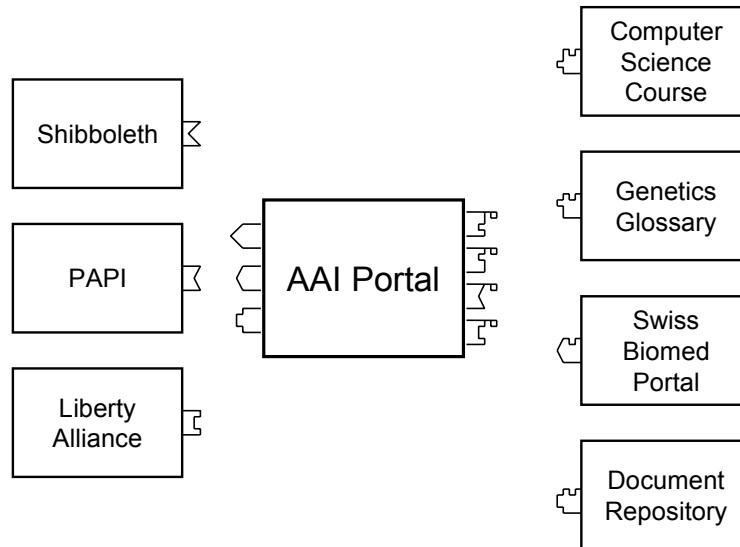


Figure 3: The AAI portal with its interfaces to AAI and resources.

The AAI portal represents an AAI-enabled resource towards the authentication and authorization infrastructure, as shown in Figure 4. Resources that are plugged into the portal are AAI portal-enabled and get their user information attributes from the AAI portal.

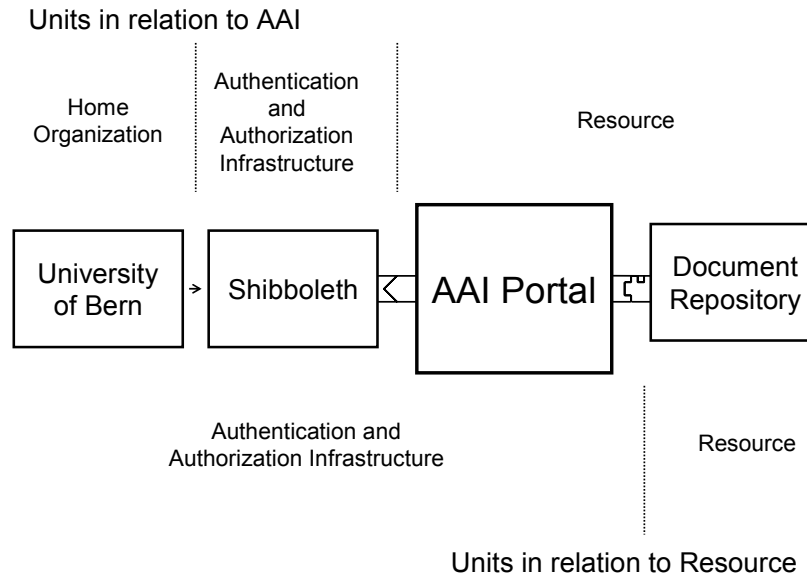


Figure 4: Units in relation to the AAI portal.

AAI portal administrators can add new resources by selecting among the built-in resource adaptors such as described below. Each built-in resource adaptor provides information on its functionality and on how to integrate the resource. The integration of resources is supported by resource adaptor provided forms where the required information, such as URL or shared secret must be entered. Since the AAI portal prototype implementation has been written in PHP, many resource owners are able to develop or modify resource adaptors. The integration into the AAI portal is facilitated by a detailed explanation of the integration procedure.

The AAI portal does basically the same as the AAI; it redirects users via HTTP from the AAI portal to resources. The AAI portal in most cases has to release user information to the resource. The procedure that defines how users take their information from the AAI portal to the resource is delegated to the resource adaptors. Many procedures are possible, for example cookie-based solutions or based on GET/POST methods in the URL. User data can also be written into the resource's database as implemented with the Institutional Management System (IMS) [9] API. In most cases, a shared secret between the AAI portal and the resource has to be established and the exchanged data encrypted.

One of the already implemented resource adaptors uses the IMS API and is used to redirect AAI portal users to a course home page hosted on a WebCT 4.x and Vista x server [10]. The WebCT adaptor passes authentication data directly to the WebCT user database in order that the user is automatically signed on. When a user tries to access the WebCT course guarded by this adaptor, the adaptor first creates a WebCT user for the given AAI user if it does not already exist. It also subscribes the user to the course. This is carried out by invoking the respective calls in the WebCT IMS API. Secondly, the resource adaptor redirects users to the course home page via the WebCT Autosignon API.

5 User and Resource Management

Resources that are directly connected to an authentication and authorization infrastructure must include user management functions. In many resources, only a very basic or even no user management exists. For example, still many resources consist of HTML pages, sometimes .htaccess protected.

Each user accessing the AAI portal can act in different roles. We propose to implement at least three different user roles. The super user is the AAI portal administrator and allowed to configure everything on the AAI portal. The second user role is the resource administrator. He acts as owner of resources hosted on the AAI portal. The third user role is the resource user, sometimes simply called 'user'. These are persons accessing resources hosted on the AAI portal.

The AAI portal provides basic user management functions. An additional advantage is that users can directly be added to the portal as local users if needed. These users access the AAI portal through an access interface on the portal itself and not through the AAI. The local access is also used by portal administrators for at least the initialization procedure of the portal. The local user access gives the possibility to temporarily add students to the resource without going through the process for adding them to a home organization.

Figure 5 shows the origin and flow of user information attributes. Home organizations provide user information attributes to resources based on their attribute release policy; this can further be restricted by each user. Resources define their attribute acceptance policy and grant access only if potential users conform to this policy. This process has been described in the example of Alice and the TCP/IP course. In the case a home organization cannot provide a certain user information attribute or a resource owner would like to know something about its resource users that is never stored in a home organization's database; users can provide personal information directly to the AAI portal. This is a very important feature provided by the portal because home organizations may probably only release very basic user information attributes. Since user provided information is not reliable, compared to the home organization provided, it is represented to administrators as user provided data. User provided data is about as reliable as Email addresses or mobile phone numbers, which are collected during lectures. The resource owner of the TCP/IP course would like to know if students have already read literature about the Internet Protocol. This information is essential to him for the preparation of supplementary readings for the new students. It is very unlikely that a home organization ever will store a user information attribute called "IpKnowledge" for example. This attribute can be self-defined and added to the existing attributes on the AAI portal and all TCP/IP course subscribers have to provide this information.

Figure 5 also shows that user information attributes can optionally be provided by the AAI portal to resources, but there is no requirement for this action. It only depends upon the chosen method a resource must receive information and of the way the resource adaptor performs the user redirection.

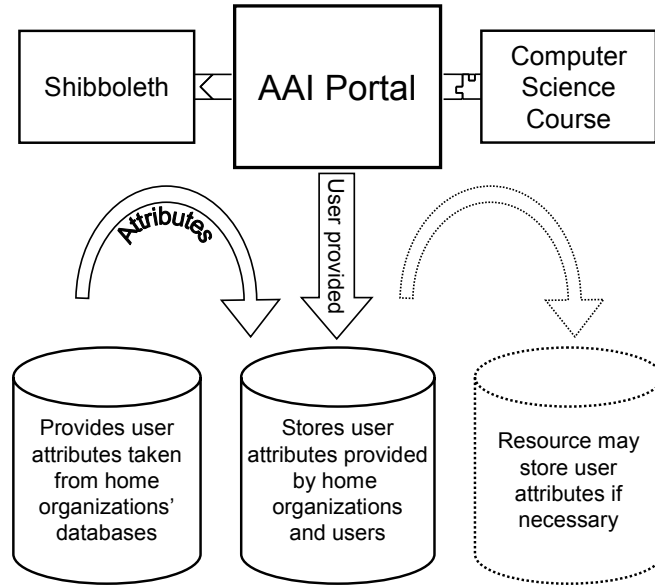


Figure 5: Origin and flow of user information attributes.

User information stored in the AAI portal is released to resources based upon the user's attribute release policy. Users define an attribute release policy for each AAI portal resource they subscribe to. This guarantees that different AAI portal resources on the same portal do not automatically get the entire user information.

Resources can be open or closed for subscription and open, closed or suspended for subscribers. A resource can grant access to all subscription requests or add subscription requests to a waiting list. Resource owners define their attribute acceptance policy by selecting predefined AAI attributes or creating custom ones. Resource administrators can list their respective subscribers and view the information they have provided.

Resource administrators have the possibility to inform users about user and resource status changes by Email or short message service (SMS).

6 Value Added Services

The AAI portal represents a point of access to all its hosted resources. Users re-access the portal each time they access an AAI portal hosted resource. The AAI portal's architecture foresees the integration of value added services, which can be plugged into the portal upon demand. Those specially adapted resources can perform services not integrated in the target resource. The resource owner thus outsources certain services to the AAI portal. The services are integrated into the AAI portal through an AAI portal administrator and can be configured by resource administrators.

The biggest advantage of these value added services lies in their personalization. Resource administrators are both owners and administrators of their services. Users are automatically subscribed to the value added services that belong to their subscribed resources.

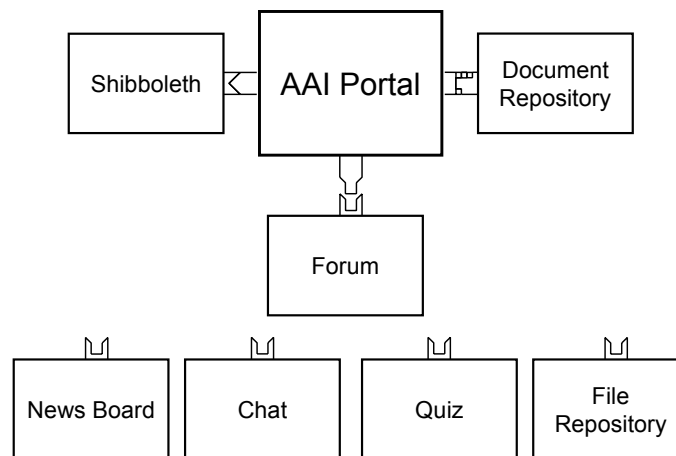


Figure 6: The AAI portal and a selection of pluggable value added services.

Figure 6 shows the AAI portal, here connected to Shibboleth as authentication and authorization infrastructure together with a document repository as connected resource. The resource owner would like to add a forum to his resource without changing the resource platform that does not foresee a forum. He asks the AAI portal administrator to add a forum to the portal and the forum is ready to use. In the existing portal prototype implementation, a discussion board, a document repository and a chat have been integrated in the above described way.

7 Discussion and Conclusions

All concepts described above have been implemented and released. The release consists of PHP code and uses a MySQL database. The broker portal has been connected to the Swiss Shibboleth implementation [11] operated by SWITCH. SWITCH has taken over the AAI portal's project lead and maintenance of the code, which is licensed under the general public license. The AAI portal's home is now on the SourceForge web site. A set of resource adaptors has been implemented. It covers a big palette of resources such as cookie and password protected web sites, WebCT CE and Vista course platforms and more.

University of Bern uses the broker portal for connecting the course platform WebCT CE to the Swiss-AAI. On the same portal hosted is a genetics glossary with many national and international students and the above described computer networks laboratory VITELS. Other portals are operational for a nano science laboratory [12], a biomedical portal and Edutech. Edutech is currently (2004) setting up a portal for connecting the course platform WebCT Vista, which will serve the Swiss Virtual Campus [13] e-learning projects.

The concept of filling the gap between authentication and authorization infrastructures and resources by providing user and resource management functions and enhancing user information collection proved its usefulness.

8 Outlook

The integration of additional value added services will make the AAI portal much more interesting for content providers that do not own an e-learning course system. Enhancing the AAI portal with additional accounting functions will ease its use in pay per use applications. Especially, a provision for financial accounting is required. Developments in this direction have already begun. Financial accounting will be very interesting for resource providers, especially because it is currently not addressed in AAI. We also think of integrating content adaptation systems, which adapt media content to user profiles, for example to mobile or broadband users.

9 References

- 1 Steinemann M.-A. et al. (2002). Global Architecture and Partial Prototype Implementation for Enhanced Remote Courses, Computers and Advanced Technology in Education (CATE 2002), Cancun, Mexico, ISBN 0-88986-332-6, pp. 441-446
- 2 SWITCH. The Swiss Education & Research Network. <http://www.switch.ch>
- 3 SourceForge. <http://www.sourceforge.net> and <http://aai-portal.sf.net>
- 4 Cantor S., Erdos M. (2002). Shibboleth-Architecture DRAFT v05. NSF Middleware Initiative Draft.
- 5 RedIRIS. Spanish National Research Network. <http://www.rediris.es>
- 6 Castro-Rojo R., López D.R. (2001). The PAPI System: Point of Access to Providers of Information. Terena.
- 7 Wason Tom (2003). Liberty ID-FF Implementation Guidelines. Draft Version 1.2-02. Liberty Alliance Project.
- 8 Microsoft .NET Passport. <http://www.passport.com>
- 9 IMS Global Learning Consortium, Inc. <http://www.imsproject.org>
- 10 WebCT Web Course Tool. <http://www.webct.com>
- 11 Graf Christoph et al. (2003). Architecture Evaluation. SWITCH. <http://www.switch.ch/aai>
- 12 Guggisberg M. et al. (2001). An Interdisciplinary Virtual Laboratory on Nanoscience. Electronic Notes in Future Generation Computer Systems, Elsevier, Vol. 1
- 13 Swiss Virtual Campus. <http://www.swissvirtualcampus.ch>
- 14 SWITCH AAI. <http://www.switch.ch/aai>

10 Acknowledgments

We would like to acknowledge the contributions to the work presented in this article of all involved persons. We especially appreciate the financial support by the Swiss Virtual Campus and the technical support by the SWITCH AAI team [14].