bwIDM

Federating IT-based Services at Baden-Württemberg's Universities

In Baden-Württemberg, universities provide a wide range of IT-based services like high performance computing, data storage and analytics facilities, to name a few. Various services are operated only at one or a few places, but should also be accessible from other sites. How can federated identity management be leveraged to facilitate state-wide secure access?

Published by the project management
of the bwIDM project, represented by:

Karlsruhe Institute of Technology (KIT)
Steinbuch Centre for Computing (SCC)
Zirkel 2, 76131 Karlsruhe, Germany

Contact: Dr. Martin Nußbaumer
phone: +49 721 4 608 8073
martin.nussbaumer@kit.edu

December 2013 (SCC-TB-2013-1)
Editor: Dr. Sebastian Labitzke

http://www.bwidm.de

A Project Funded by the Ministry of Science, Research,
and the Arts of the State of Baden-Württemberg, Germany.

Prof. Dr. Hannes Hartenstein
Project Manager

The *bwIDM* project has been a joint effort of the information technology centers of the universities of the state of Baden-Württemberg to establish an Authentication and Authorisation Infrastructure (AAI) that meets the specific needs for state-wide IT service provisioning and usage in academic environments.

An AAI is an essential element for IT security as it makes sure that services can only be accessed by users that are authorized for using this service. A *good* AAI, thus, makes access easy for authorized users – and very hard for all others.

Why was there a need for a "Baden-Württemberg solution"? There are two main reasons: first of all, various non web-based services like resources for high performance computing and data management had to be federated. As a solution for this challenge was not available, we had to come up with one. Second, as various IT services are funded only for the use by members of the universities in Baden-Württemberg, we were in need for a sub-federation to the nationwide AAI provided by DFN e.V.

We hope that the *bwIDM* approach is not only considered as a solution for Baden-Württemberg, but also considered as a blueprint for others – a solution "made in Baden-Württemberg".

I would like to thank the whole *bwIDM* project team for their enthusiastic and creative work. In particular, I would like to thank Martin Nußbaumer for great project leadership. He together with Michael Längle, Thomas Nau and Saher Semaan led the core team – a big thanks to all of you! All universities in Baden-Württemberg participated in the project, members of the computing centers of the universities of Freiburg, Konstanz, Ulm, and the KIT formed the core team – a big thanks also to Gerhard Schneider, Marcel Waldvogel, Hans-Peter Grossmann and Stefan Wesner for their backing. Last, but definitely not least, a big thank you to Peter Castellaz, MWK.

It was an honor and a pleasure to be responsible for this project effort,

Prof. Dr. Hannes Hartenstein

Content Overview

bwIDM

1 - Project Charter and Achievements

Federated identity management yields many advantages, such as enhanced usability and improved quality of identity information.

Web-based services are already successfully and widely federated using the Security Assertion Markup Language (SAML). In terms of usability and quality of identity information, non web-based services benefit from being federated similarly than web-based services. However, IT services provided within the state of Baden-Württemberg, Germany are often already deployed and in use. Although federated computing environments are also existent, e.g., the *bwGRiD* approach (www.bwgrid.de), due to usability reasons, often only predominantly computing experts are able to access the resources. Hence, more usable and convenient ways are researched to allow easy access for a broader range of users.

Up to this point, no versatile approach has emerged to federate non web-based services that can easily be integrated. Consequently, the project *bwIDM* dealt with technical and non-technical challenges in the context of building up *trust* between identity providers in the state of Baden-Württemberg with the goal to federate IT services – web-based and, in particular, non web-based.

The contribution of *bwIDM* is the establishment of this trust in form of the *bwIDM* federation with appropriate federation technologies as the foundation for federating state-wide IT services. Thereby, two main objectives concerning usability and deployability were addressed by the project:

- **The easy access to state-wide IT services for researchers in Baden-Württemberg**: in the state of Baden-Württemberg, researchers can access decentralized web-based and non web-based services by the use of their local account.

- **The easy entry to the *bwIDM* federation for new and existing IT services**: the researched and deployed federation technology aims at minimizing the integration effort in terms of both usability and necessary adjustments to existing or future service deployments.

Since it was a critical success factor to keep the local identity management systems independent of the technical requirements of *bwIDM*, only a set of rules and interfaces were adopted to regulate the interaction. The following sections illustrate the overall results and findings of the *bwIDM* project.

Dr. Martin Nußbaumer

2 - Federating non Web-based Services via an LDAP Facade

One of the major deliverables of the *bwIDM* project is the developed concept [1][3][4] (see Section 5) to federate non web-based services via the Security Assertion Markup Language (SAML), i.e., services that are not accessed via a web-browser. As SAML Identity Providers (IdPs) are already established at the universities and most Service Providers (SPs) rely on proprietary software that is hard to modify, an easy integration of the *bwIDM* concept both with existing IdPs and SPs is of paramount importance. Furthermore, in case of SSH services, such as compute clusters, most users already have an of-the-shelf SSH client preinstalled that should be compatible with the *bwIDM* concept without any modifications. Other important issues that have to be considered are legal aspects. For instance, to adhere to privacy laws, an IdP has to inform the user of the transfer of personal data to the SP. Furthermore, SPs might require the user to consent to acceptable-use policies before granting service access.

To address those legal aspects, the *bwIDM* approach makes use of the rich client interface of conventional web-browsers by requiring the user to register via a registration web application. Since authenticating to this web application is the primary use-case for SAML, existing plugins like uApprove (https://www.aai.dfn.de/dokumentation/identity-provider/konfiguration/uapprove/) can be utilized by the IdP to inform the user of the transfer of personal data to the SP. Furthermore, once the user is logged in at the SP, the consent to acceptable-use policies can be requested and a local account is provisioned that contains information that is not maintained by the IdP, such as the user's home-directory.

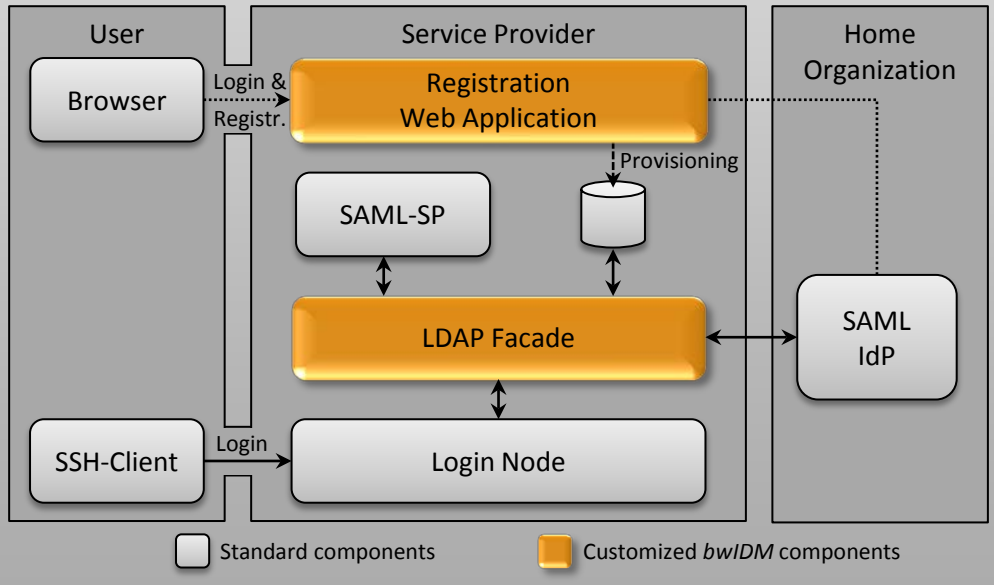

Figure 2: Architecture of the bwIDM solution (Option: LDAP Facade)

However, this local account is not sufficient to authenticate and authorize a user when accessing the actual service, as (1) the user should be able to authenticate via the credentials of the home-organization and (2) the authorization decision should not be based on potentially stale authorization attributes. That is, a successfully registered user should not be able to access the service if he/she lost the necessary access requirement at the home-organization at a later point in time. Thus, each time a user accesses an SP, the SP has to query the IdP for up-to-date attributes.

To enable (1) and (2) without requiring any changes to the services themselves, the *bwIDM* approach leverages standard interfaces such as the Lightweight Directory Access Protocol (LDAP) (cf. Figure 2) or Plugable Authentication Modules (PAM). While the Facius approach [3] can be used to integrate unmodified services that offer PAM support, the LDAP-Facade approach [4] can be used to "hide" the federative credential and attribute checking behind an LDAP interface that can be used by the services like a regular LDAP server.

The advantages of this approach are that no service has to be adapted in any way to be federated, IdP implementations do have to be adapted and regular, unmodified user clients can be used to access the service.

However, unmodified service clients imply that the home-organizational password of the user is sent to the SP that, in turn, verifies it against the IdP. In case the SP and IdP do not belong to the same organization, this can conflict with IT security policies. To address this problem, the *bwIDM* approach allows users to establish alternative authentication methods during the registration process such as an SP-specific password or a deployed SSH-PublicKey. Alternatively, if the service client can be modified, the user password can be sent directly from the client to the IdP, which then issues assertions that can be passed to the SP to prove a successful authentication. In particular with services, such as the *bwSync&Share* service (see Section 4.1), this solution is preferable, as users have to download a new client anyway.

The developed concept is already actively used to federate access to *bwHPC* services (see Section 4.4), such as *bwUniCluster*, and *bwLSDF* services, such as *bwSync&Share* and *bwFileStorage*. However, the concept is by no means limited to the *bwIDM* use-case and can be applied to federate arbitrary non web-based services via SAML.

Jens Köhler, Michael Simon

3 - Federation Management and Interoperability with the DFN-AAI

Another main goal of the *bwIDM* project was to establish a **non-isolated** and **integrated** federated Identity Management System (IdM-System) for the universities of the state of Baden-Württemberg within **already existing and well established** workflows and structures. A major challenge was to create and establish an environment to be able to exchange user information within the federation in a compliant manner (according to the Federal Data Protection Act). The **German National Research and Education Network** (DFN e.V.) provides a well established **Authentication and Authorization Infrastructure** (**DFN-AAI,** https://www.aai.dfn.de/en/) for the universities of Germany, which coordinates the interoperability between national and international research institutions. Within the DFN-AAI sub-groups can be formed that consist of members agreeing on specific details to extend the scope of interoperability and cooperation.

In cooperation with the DFN-AAI, we succeeded to integrate the *bwIDM* solution into the national AAI by also relying on SAML as backend technology. Additionally, this integration allows to participate (almost automatically) in international AAIs, such as **eduGAIN** (http://www. geant.net/service/eduGAIN/) as a Pan-European Web Single Sign On (Web SSO) service.

The technical integration of the *bwIDM* project into the existing DFN-AAI infrastructure is realized by extending a metadata administration tool implemented and operated by the DFN-AAI. This extension allows an arbitrary set of attribute information to be carried within an entity's (or a group of entities') metadata to communicate additional information about that entity (or group) to a metadata consumer. This mechanism allows the operator (in this case the DFN-AAI) to include additional information about the sub-federation's member sites (e.g., their adherence to optional federation policies) signed by the *bwIDM* committee.

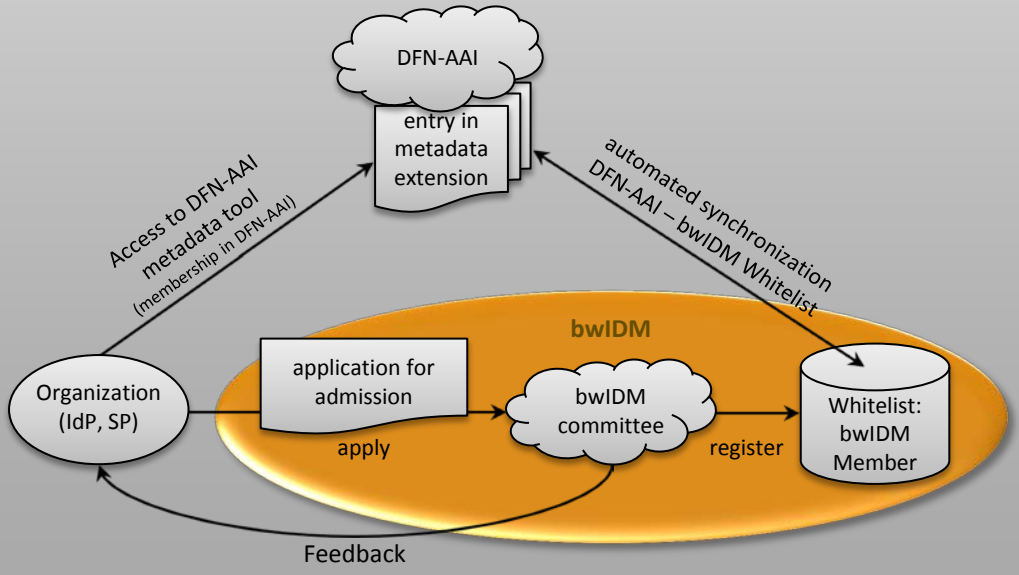

Figure 3: Interoperation of bwIDM and the DFN-AAI

Metadata consumers can then choose to process or ignore such information as they deem necessary (cf. http://docs.oasis-open.org/security/saml/Post2.0/sstc-metadata-attr.html).

In other words: with the previously mentioned DFN-AAI metadata administration tool a *bwIDM* affiliated Identity Provider (IdP) or Service Provider (SP) can be labeled with the entity category "**bwidm-member**", which is a special type of entity attribute of the aforementioned metadata extension. Relying on this additional information an SP can utilize a filter on the signed metadata of the DFN-AAI federation to grant access to specific resources for all users of the sub-federation *bwIDM*. Furthermore, a *bwIDM* affiliated IdP can release or filter additional information (attributes) about the users for *bwIDM* affiliated service providers.

Figure 3 describes the workflow of the cooperation between *bwIDM* and the DFN-AAI. In addition, it describes the procedure for applying for a *bwIDM* membership.

Saher Semaan

4 - Service Integration

The previous sections demonstrate that *bwIDM* constitutes an enabler for integrating web-based *and* non web-based IT services into an enhanced federation due to a well defined identity management infrastructure. Based on the technical *bwIDM* concepts presented in Section 2 and based on the concepts to cooperate with a national AAI (see Section 3), several IT services have already been integrated into the *bwIDM* federation. Furthermore, a significant number of IT services is planned to be integrated into the IT service federation based on the bwIDM solution.

The already integrated IT services demonstrate not only the applicability of the *bwIDM* solution in general, but also the deployability of IT services that are enhanced by the *bwIDM* concepts. Additionally, the already productive IT services show that the *bwIDM* solution can be deployed without any loss of the IT services' maintainability and operability. With bwIDM, we propose a sustainable solution that tackles the emerging issues, when web-based, as well as non web-based IT services are to be federated. As stated before, *bwIDM* constitutes not only a solution applicable for IT services provided in the state of Baden-Württemberg, but also a concept that can be used for further (sub-)federations.

In the following, we report on some of the already implemented, as well as on a selection of planned IT service integrations, which are based on the *bwIDM* solution.

Dr. Sebastian Labitzke

4.1 - bwSync&Share

A Federated Service for Synchronizing and Sharing Documents

The *bwLSDF* research project investigates the advantages and drawbacks of centralized and distributed storage systems and identifies the potentials and risks of federated storage usage. Besides traditional NAS and SAN concepts, we are looking into new approaches of flexible storage management. Our objective is the development of a new service that offers access to the Large Scale Data Facility (LSDF) for all students and employees studying and working at universities in the state of Baden-Württemberg.

The exchange of documents via USB-drives or e-mail is limited and no longer practical. In January 2014, we introduce a service called *bwSync&Share*, which is a privacy-aware alternative to the well-known dropbox (https://www.dropbox.com/). The new service allows documents to be synchronized and shared between different users and devices. The service offers synchronization clients for the most common platforms (Windows, Linux, MacOS, iOS and Android), as well as an intuitive web interface to access and share personal documents. The users benefit from easy-to-use, collaborative and platform independent workflows, which enrich the scientific and academic processes (e.g., writing scientific papers or research proposals).

During the deployment of the *bwSync&Share* service, we faced different challenges. The most important issue was the authentication and authorization of users in a federated, state-wide context. Since the service is available to about 80 organizations and approximately 450,000 users, a federated Authentication and Authorization Infrastructure (AAI) is essential.

The *bwSync&Share* service is built on top of an AAI that was drafted, designed and implemented within the *bwIDM* project. The infrastructure is based on a system called Shibboleth that implements the Security Assertion Markup Language (SAML) standard to forward authentication and authorization information of users from an Identity Provider (IdP) to a Service Provider (SP). The pre-established trust between IdP and SP enables a federated identity management where the IdP is in charge of providing verified user identities and the service provider accepts the transmitted credentials. This functionality fulfills the *bwSync&Share* requirements perfectly and, therefore, was chosen to be used.

In the context of providing state-wide IT services, a federated authentication and authorization infrastructure is a fundamental requirement that applies not only to Sync&Share solutions. The *bwIDM* team faced this challenging task and mastered it splendidly. The project's outcome is a solid foundation, which enables all subsequent federated IT services in the state of Baden-Württemberg.

Nico Schlitter

4.2 - bwHPC

Federated Identity Management for High Performance Computing in the State of Baden-Württemberg

High performance computing in the state of Baden-Württemberg has a long tradition. In the current bwHPC concept, the HPC infrastructure covers all tiers and correspondingly varies with respect to peak performance and scalability (Figure 4).

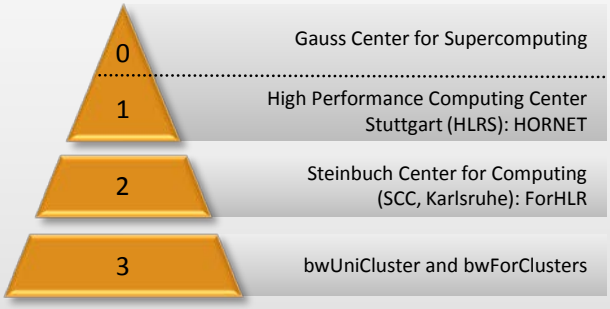

Figure 4: High performance computing in Baden-Württemberg

Tier 3 is commonly referred to as the HPC entrance level. HPC systems of tier 3 represent local resources of research facilities and universities. Providing a common HPC entrance level for users, the state of Baden-Württemberg started to coordinate activities of user support and administration services amongst tier 3 installations and, therefore, initialized federated HPC at tier 3. However, federated systems require identity management across all involved organizations and locations.

Federated HPC in Baden-Württemberg first emerged as *bwGRiD* in 2008. Co-financed by D-Grid and embedded in its service infrastructure, *bwGRiD* made use of a cryptographic certificate system for its identity management. While certificate based systems have their advantages for very large scaled grid computing infrastructures, the capabilities for smaller scaled and heterogeneous HPC infrastructures are limited. Users of *bwGRiD* are faced with a complicated multi-step registration procedure including issuing a personal certificate and the membership of a virtual organization before access is granted to the *bwGRiD*. Moreover, personal certificates need to be periodically reissued and managed together with a middleware in order to allow web portal or console access to HPC resources. As a result, many potential users hesitated using *bwGRiD*.

Furthermore, *bwGRiD* cluster providers faced out-of-date identity information caused by the large latency between its certificate based identity management and identity management of users' home-organizations.

In 2013, Baden-Württemberg has started the installation of a new heterogeneous infrastructure of clusters for tier 3. These clusters will be science domain specific and constitute the new federated HPC. Based on the success of the project *bwIDM*, a federated identity management has already been implemented and deployed for the bwUniCluster coupling home-organizational and HPC local accounts. Registration and authentication processes via the users' approved home-organizational accounts will ease the access. Finally, *bwHPC* tier 3 will become a true HPC enabler.

Dr. Robert Barthel, Tobias König

4.3 - bwHPC Science Portal
bwIDM-based Cluster Access

The *bwGRiD/bwHPC science portal* is a GridSphere-/Liferay-based web server integrating tools for scientific research, cluster access and teamwork functionality below a secure unified access point. Users can manage jobs at multiple clusters through an easy-to-use web interface. Portlets provide interfaces to applications in the areas chemistry, mathematics, biology, computer-aided engineering, medicine and physics. Some of the portlets offer graphical user interfaces to create input files used by the applications. General purpose portlets support job monitoring and data management. Scientific research groups and teaching classes can exchange input data sets or results by means of teamwork repositories. An Avatar provides detailed help for most input elements of the portlets, can guide new users step-by-step through specific tasks, and can respond to frequently asked questions.

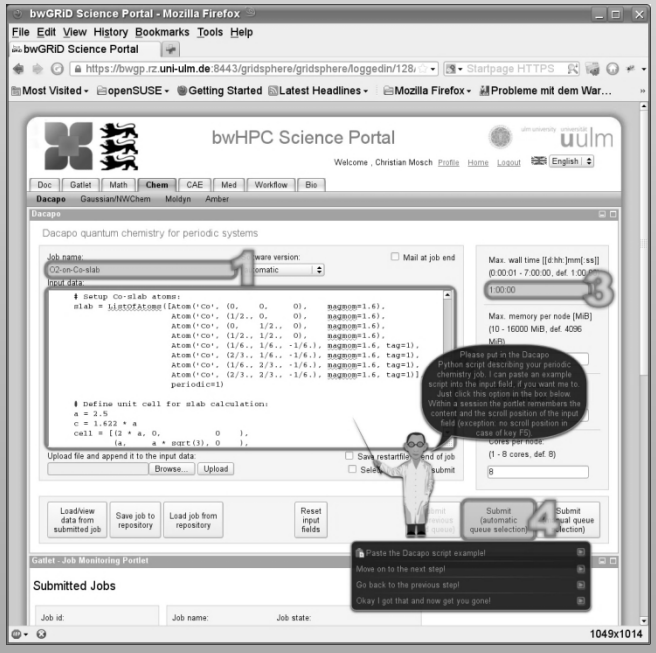

Figure 5: bwHPC Science Portal

Since October 2011, the science portal is available at "https://portal.bw-grid.de/". So far, more than 200 users utilized the portal to access the *bwGRiD* clusters. A survey showed that new users are significantly faster when accessing the clusters and their applications through the portal. With instructions on how to use the command line, a given task could be accomplished in roughly 54 minutes. By means of the portal, but without the Avatar, the task takes 22 minutes. Finally, with the help of the Avatar, the task can be fulfilled in only 12 minutes. So the portal has the ability to get new users much faster to productivity.

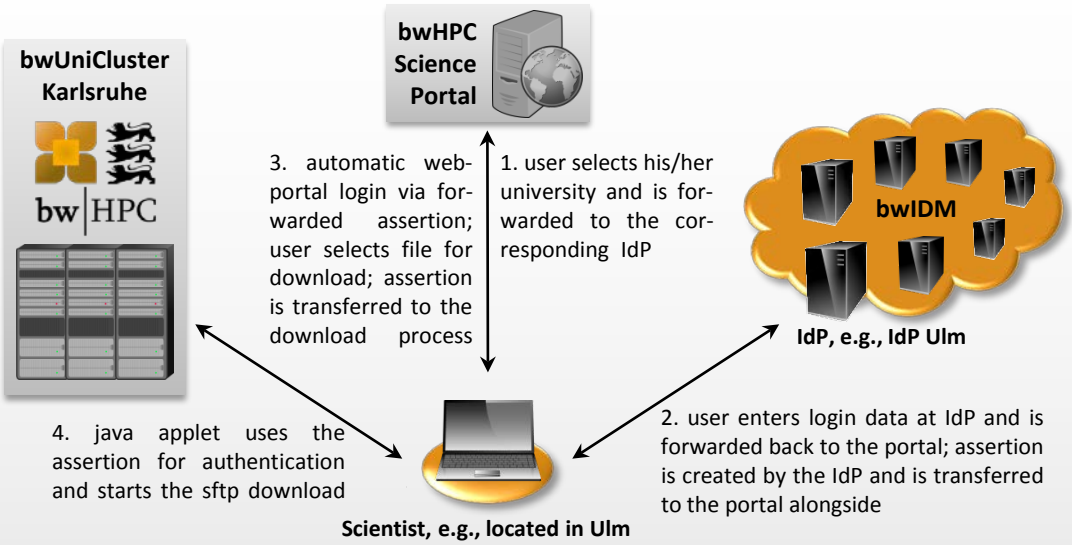

Figure 6: Workflow of the bwHPC access via the bwHPC Science Portal

The new Liferay-based version of the portal, which makes use of Shibboleth authentication, replaced the time-consuming registration procedure for grid certificates. As part of the login process, the user is forwarded to the identity provider (IdP) of his/her home-organization and has to enter the login credentials of his/her local account. If the credentials are valid, the IdP creates a singed authentication assertion and forwards this assertion back to the portal. The user's Shibboleth-based portal account is then created automatically upon first login. Although the user still has to register once for the compute services he/she intends to use, even this registration process is Shibboleth-based and, therefore, it can be accomplished within a few minutes. Finally, with the aid of a new delegation mechanism based on SAML assertions developed by the *bwIDM* team, the user can immediately access the *bwHPC* clusters. Within this mechanism the assertion is used like a password to identify and authenticate the user at other service providers, e.g., a remote cluster. Figure 6 illustrates the login procedure and an exemplary file transfer to the user's computer.

Christian Mosch, Bastian Boegel

4.4 - bwLehrpool
Infrastructure for Virtual Laboratories

bwLehrpool has the aim to develop and construct a centralized infrastructure for universities in the state of Baden-Württemberg. These universities are capable of mutually developing and using virtual laboratories. Each laboratory can be quickly and easily used within the universities' own PC pools or specialized labs. Despite the centralization it is possible to deploy university specific characteristics by combining double abstraction and virtualization technology. These specific characteristics are, for example, authentication, private home drive, printers or license servers. It is possible for students to work with any virtual laboratory within the lecture or in their free time on any computer on site (see Figure 7).

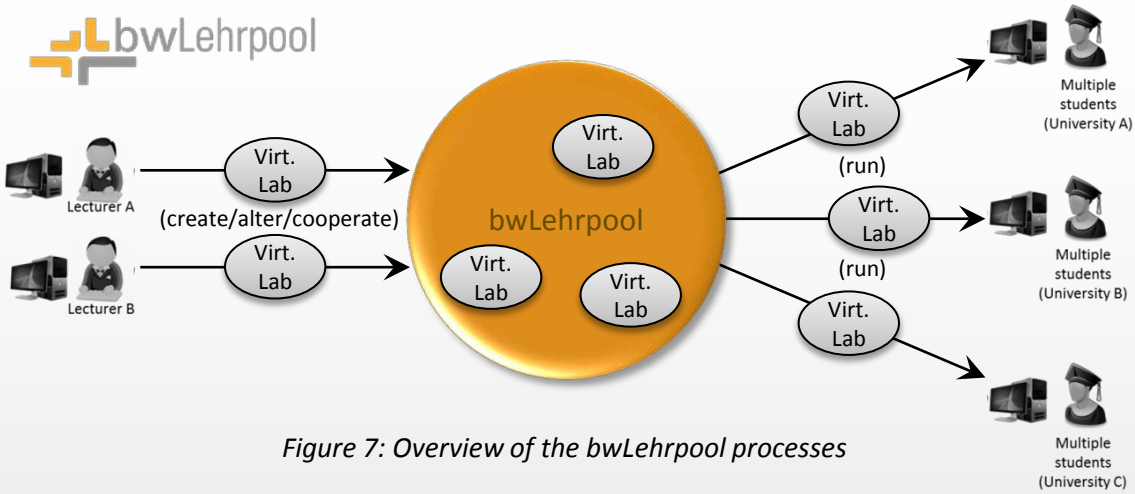

Figure 7: Overview of the bwLehrpool processes

From an administrators point of view, the main advantages are found in the reduction of administration effort, strengthening the university-wide synergies, reduction of IT costs and optimization of teaching. This way, not only the administrative effort in the individual departments can be minimized, but also the university-wide expertise on operating systems, applications, licensing and complex application environments can be used more effective. Furthermore, the usage of virtualization ensures the independence regarding hardware and lecture rooms. This means that courses and laboratories are considerably more flexible to plan and one can react to changes far more quickly. Additionally, costs can be reduced in the computer rooms due to the usage of heterogeneous hardware.

In terms of teaching, there are significant advantages as well. With the help of the intuitive lecturer module, any lecturer of any faculty can quickly develop or alter a virtual lab and provide it to any lecture room or special laboratory. The additional benefit is immense, as they can independently create a lab at any time of day with hardly any effort: from a single application to a complex system of data bases and web servers. During this process the lecturers are supported by already existing templates (an operating system including the most important configurations).

Another big advantage is the possibility of cooperatively developing a virtual laboratory. This can be done by university-wide access for multiple lecturers to develop and optimize these laboratories.

Only thanks to *bwIDM* and its state-wide authentication, the idea of centralization and university-wide cooperation can easily be set up for the future service *bwLehrpool*. In the optimal case, just based on the assignment of the role “lecturer” for lecturers of the universities, which usually happens automatically, the rollout of *bwLehrpool* using *bwIDM* can be completed very quickly.

For more information visit http://bwlehrpool.hs-offenburg.de/.

Michael Wilson, Prof. Dr. Jan Münchenberg,
Wolfgang Honigberger, Dr. Dirk von Suchodoletz

5 - Selected Publications

[1] M. Simon, M. Waldvogel, S. Schober, S. Semaan, M. Nussbaumer et al., **bwIDM: Föderieren auch nicht-webbasierter Dienste auf Basis von SAML**, Lecture Notes in Informatics (LNI - Proceedings, GI-Edition), 5. DFN Forum Communication Technologies, Regensburg, Germany, May 2012

[2] Saher Semaan, Richard Zahoransky: **bwIDM: Anbindung nicht-webbasierter IT-Infrastrukturen an eine SAML/Shibboleth-Föderation**, 8th Joint BFG/bwGRiD Conference & Workshop on High-Performance Grids, Freiburg, Germany, May 2012

[3] J. Köhler, S. Labitzke, M. Simon, M. Nussbaumer, H. Hartenstein, **FACIUS: An Easy-to-Deploy SAML-based Approach to Federate Non Web-Based Services**, 11th IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom-2012), Liverpool, UK, June 2012

[4] J. Köhler, M. Simon, M. Nussbaumer, H. Hartenstein, **Federating HPC access via SAML: Towards a plug-and-play solution**, International Supercomputing Conference, Leipzig, Germany, June 2013

6 - Project Participants

We would like to thank the Ministry of Science, Research and the Arts of the State of Baden-Württemberg for the funding, as well as the ZENDAS (Zentrale Datenschutzstelle der baden-württembergischen Universitäten) and the ALWR (Arbeitskreis der Leiter wissenschaftlicher Rechenzentren in Baden-Württemberg) for the project support. Furthermore, we like to thank all members and participants of the project bwIDM for their invaluable contributions. In particular, we like to acknowledge the following contributors :

University of Konstanz:
Michael Längle, Markus Grandpre, Daniel Scharon, Jakob Becker, Peter Ulber
University of Ulm:
Thomas Nau, Claudia Pauli, Daniel Baur, Sven Schober, Vladimir Nikolov
University of Freiburg:
Saher Semaan, Richard Zahoransky, Bernd Oberknapp
Karlsruhe Institute of Technology (KIT):
Michael Simon, Jens Köhler, Tobias Dussa, Sebastian Labitzke, Martin Nussbaumer
University of Mannheim: Tobias Kienzle, Steffen Hau, Heinz Kredel
University of Heidelberg: Reinhard Mayer, Markus Skowronek
University of Tübingen: Stefan König, Kurt Spanier
University of Hohenheim: Steffen Bücheler, Henning Reineke, Björn Breiner
University of Stuttgart: Dominik Lamp, Björn Eich
IUK-BW: Herbert Röbke

In Baden-Württemberg, Germany, universities provide a wide range of IT-based services like high performance computing, data storage and analytics facilities, to name a few. Various services are operated only at one or a few places, but should also be accessible from other sites. How can federated identity management be leveraged to facilitate state-wide secure access? In this flyer, we report on the project bwIDM. Within bwIDM a state-wide identity management system was designed and implemented to provide access to web-based and, in particular, non web-based IT services. It is illustrated how this state-wide infrastructure is integrated into a national authentication and authorization infrastructure. We demonstrate that the bwIDM solution can be easily deployed to federate arbitrary IT services. Additionally, we present some of the successfully integrated IT services, which are based on the bwIDM solution.

bwIDM – a solution "made in Baden-Württemberg".

bwIDM