

Strategy and Tools for Identity Management and its Process Integration in the Munich Scientific Network

Wolfgang Hommel¹, Silvia Knittl², Daniel Pluta³

¹Leibniz Supercomputing Centre, Boltzmannstr. 1, D-85748 Garching, hommel@lrz.de,

²Munich Network Management Team, Boltzmannstr. 1, D-85748 Garching, knittl@mm-team.org,

³Technische Universität München, Boltzmannstr. 3, D-85748 Garching, pluta@tum.de

Keywords

AAI, Identity and Access Management, Administration, Directory Services, LDAP, Group Management

1. EXECUTIVE SUMMARY

For many higher education institutions (HEIs), identity and access management (IAM) has proven to be a key enabling technology which does not only automate the handling of user accounts and their privileges to a large extent, but also allows a very tight integration into the existing business processes and offers comprehensive interfaces to IT service management processes. The Munich Scientific Network (Münchner Wissenschaftsnetz, MWN) spans multiple HEIs, including the Munich universities, and the Leibniz Supercomputing Centre as their common IT service provider. In this article, we discuss the current and medium-term strategy for closely coupling the IAM infrastructure components existing at each institution, as well as the integration of external users from both, Shibboleth federations and European Grid computing projects. For handling over 100,000 users and a few tens of services, adequate tools for delegated administration are necessary. We selectively sketch the dedicated tools, which we are developing for the decentralized management of guest accounts and group management, as well as their underlying concepts and the rationale for not using existing tools; they are intended to be customizable for other HEIs and will be made available.

The Leibniz Supercomputing Centre (LRZ) is the common IT service provider of the Higher Education Institutions (HEIs) in the MWN. It offers the full spectrum of HEI computing center services, such as groupware, file servers, backup and archiving, WLAN and VPN access, and operates the network infrastructure connecting more than 400 buildings and 60,000 end systems. Furthermore, the LRZ is one of Germany's academic supercomputing centers and operates high performance computers which are used by major German research projects and within European Grid computing projects.

Given the comparatively high number of students and employees at the Munich universities, namely Ludwig-Maximilians-Universität (LMU) Munich and Technische Universität München (TUM), the increasing number of alumni services, and the increasing number of users from the German Shibboleth federation DFN-AAI (see DFN-AAI website (2008)) as well as the European Grid projects, more than 100,000 users and their access rights need to be managed. This large number of users has necessitated the planning of identity management systems (IMs) and related tools in the involved institutions.

Meanwhile, both Munich universities as well as the LRZ have put their IMs into production, and while there are still ongoing efforts to further integrate services which do not make full use of the IM yet, there is a clear demand for a tight coupling of the IMs across institutional borders. For several of them, the functionality offered by federation technologies, such as Shibboleth, apparently is not sufficient. We discuss the strategy for MWN-wide identity management and its motivation in section 2, focusing on the LRZ as an identity data exchange hub. The integration of Grid and federation users and its impact on the data quality management are discussed in section 3. While delegated administration is a matter of principle for Grids and Federations, dedicated tools for the use by HEI units, such as faculties and chairs, are required to facilitate the level of data quality necessary for the automation of local authorization processes. We present the tools we are developing for the management of guests and groups as an essential part of the project IntegraTUM in section 4, where we also sketch the establishment of MWN-wide conventions for the closely related system and service administration processes. This article is concluded by a summary and an outlook to our next steps.

2. MOTIVATION AND STRATEGY FOR INTEGRATED IDENTITY MANAGEMENT IN THE MUNICH SCIENTIFIC NETWORK

Identity management across institutional borders is quite a new concept for German HEIs, although especially federation technologies, such as Shibboleth, have been in use in quite a few European countries for several years. Since the German National Research and Education Network (Deutsches Forschungsnetz, DFN) has started to operate the Shibboleth-based authentication and authorization infrastructure (DFN-AAI) in 2007, Munich HEI members can make use of an increasing number of third party services, including digital library and academic publisher services, e-learning systems at other HEIs, and the distribution of licensed software from several vendors such as Microsoft.

However, the effort required to set up and operate the required software components, i.e. the Shibboleth Identity Provider (IDP) and its underlying application, web, and sign-on servers, has turned out to be still too complex and respectively resource demanding for many HEIs themselves, for example within their departments that also run the student administration software. Therefore, both LMU and TUM have decided to have their federation software hosted at LRZ, which has taken part in German Shibboleth test beds for several years already and thus has both, an adequate server hosting environment and the necessary service knowledge. While the constellation of outsourcing the operation of a Shibboleth IDP may be a bad idea from the privacy and data protection perspective in general, the LRZ already has been contracted to process data for both universities and is also operating, e.g. TUM's IMS, so only minor extensions to the existing contractual agreements were necessary.

Sharing IT services between the HEIs in the MWN has also become an important issue. Besides cross-institutional research and collaboration projects, this primarily affects students of one of several collocated study courses, such as medicine and bio-informatics, in which the teaching of classes is teathed between LMU and TUM. Thus, students enrolled in these courses should be able to use both universities' web portals with services such as online registration for tests or grade notification, e-learning systems, and so on. This especially also concerns self service processes such as changing one's postal address: If a student changes her address in one university's self service web portal, the other university's student administration should also be informed. Unfortunately, these use cases exceed, e.g. Shibboleth's functionality: Service providers may read the student's data from the responsible Shibboleth IDP, but they cannot write or update the data stored at the IDP. Thus, current federation technologies are insufficient for such a tight integration of IMS and the underlying HEI processes; as a consequence, the traditional meta-directory approach has been chosen to solve this problem. Figure 1 shows the high-level identity management architecture for the MWN, focusing only on LMU, TUM, LRZ, and the DEISA Grid project to improve readability. The LRZ operates three LDAP-based directory services whose data is constantly being synchronized using a central meta-directory to avoid data inconsistencies:

Portal Directory Service The portal directory service is the backend of a self service portal which is used by LRZ administrators, delegated administrators in the HEIs, and users. It provides all relevant data in an LDAP schema (data format) which has been designed to support data processing in popular web portal scripting and programming languages, in our case especially Perl and PHP.

Authentication Directory Service The authentication directory service is based upon standardized LDAP schemas, including inetOrgPerson and posixAccount, and can be used by LDAP-enabled systems and services to authenticate and authorize their users without requiring their own local user database. For services which are not capable of using LDAP servers, provisioning workflows have been implemented to automate the creation, modification, and deletion of accounts in their local user databases.

Import and Export Directory The import and export directory is a meta-directory itself. It is used to retrieve identity data from LMU, TUM, and the Federation and Grid projects as discussed in section 4. It is also used to export identity data records on a very restricted selective basis. For example, account data of LMU's medicine students is not only imported at LRZ, e.g. for the email service, but also forwarded to TUM's IMS to facilitate the use of TUM's IT services within this collocated study course.

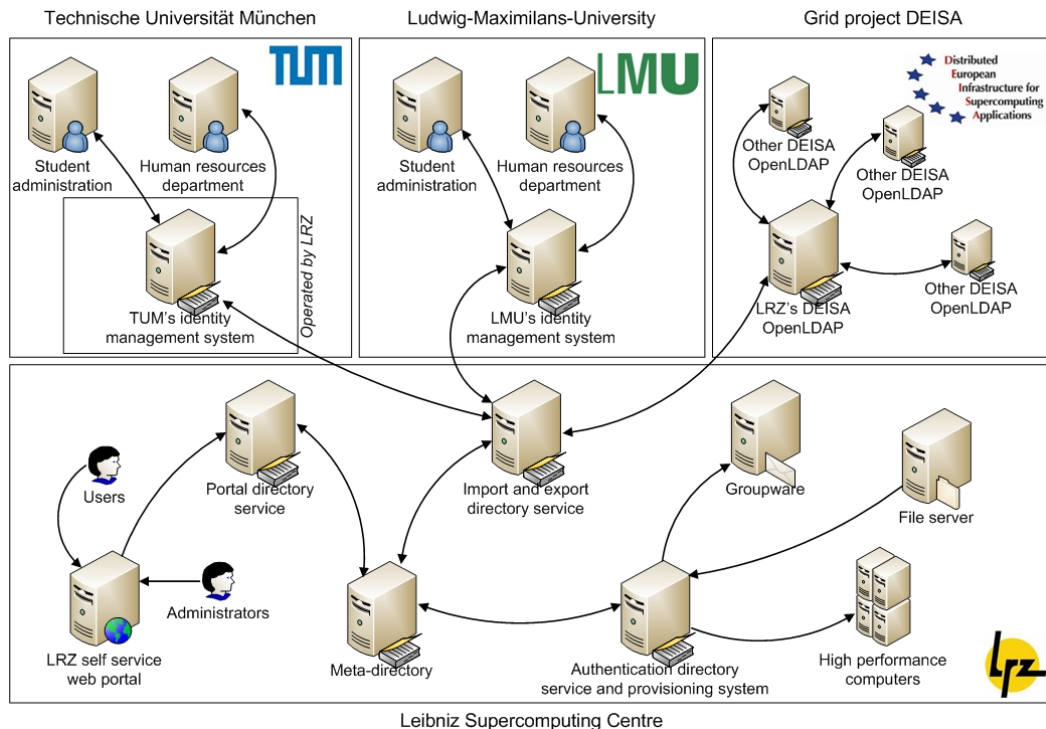


Figure 1 Coupling of identity management systems in the Munich scientific network

Summing up both aspects, the LRZ as common IT service provider does not only aggregate identity information from the HEIs, but also is intended to act as an identity data hub for the exchange of student data between them. With the additional hosting, and partially even operating of IMS components for the HEIs, the vision of identity as a service is a clear ambition in the MWN.

3. INTEGRATION OF GRID COMPUTING AND FEDERATION USERS

Both, European Grid projects and the Shibboleth-based German federation DFN-AAI, have in common that data about users from outside the MWN must be fed into the local IMS. While Shibboleth is intended for the cross-organizational use of web-based services, such as the universities' e-learning systems, Grid users require system-level access to CPU and storage capacities based on tools like secure shell (SSH) and GridFTP. Although the user data is being acquired and maintained by the responsible identity providers (in the Shibboleth case) and home sites (in the Grid case), and thus this inherent delegated administration reduces the local management overhead, the interfaces to this data are very different on the technical level. Selected user data is required locally within the MWN, e.g. for accounting and statistics purposes. As a consequence, local user profiles must be created and maintained based on the remotely available data. In the following subsections, we exemplarily describe our user management for the DEISA Grid project and the DFN-AAI Shibboleth federation. Consequences for the data quality management are discussed in subsection 3.3.

3.1. Integration of DEISA Grid Computing Users

DEISA is an EU-funded research project aiming to establish and operate a European-scale Grid computing infrastructure, which currently spans eleven high performance computing centers. Their computing and storage resources are bundled and made available to scientists for the submission of so-called 'Grid jobs' through a dedicated middleware. The use of the General Parallel File System (GPFS), a distributed file system, necessitates identically configured local accounts for each user at each DEISA site. For the management and distribution of user data, a distributed LDAP architecture has been set up:

Each DEISA site is operating an OpenLDAP service, to which those of its users must be registered which want to make use of the DEISA infrastructure. These OpenLDAP servers are connected to each other by LDAP referrals, so when data about a remote user cannot be found locally, the OpenLDAP

server at the user's home site must be queried. In our meta-directory approach, information about all DEISA users shall be retrieved and stored in LRZ's IMS. Due to DEISA's approach to user management, this poses a couple of challenges, especially the following:

Login names, as well as UNIX group names, user ids, and group ids are assigned by the user's home site. The namespaces and number ranges may clash with local accounts that have been created previously. While precautions have been taken for DEISA on an organizational level, the present solution is not scalable for the increasing number of Grid projects the LRZ takes part in. The use of referrals results in the necessity of direct LDAP access to all DEISA OpenLDAP servers for each client. As security measures are in place also on the network layer, i.e. packet filtering firewalls are used, this results in severe manual administrative overhead and additional dependencies which must be taken into account as part of the change management process.

Due to technical limitations of the meta-directory software we use, referrals to temporarily unavailable OpenLDAP servers do not result in errors. Instead, to the LRZ import and export directory service it looks like the other DEISA site has deleted all of its users. To prevent the premature deletion of accounts and the data in their home directories, grace periods must be used. However, these grace periods prevent the immediate deletion of accounts, e.g. in cases of resource abuse, which then require quick administrative intervention. Based on these experiences, the analyses of shortcomings of the current DEISA user management infrastructure as well as possible enhancements will be part of our future work.

3.2. Integration of DFN-AAI Shibboleth Federation Users

Shibboleth allows web applications to retrieve profile data from the user's identity provider. For example, e-learning systems require information about the user's name, study course, and matriculation number if passing a final test shall result in a credit the student can submit to her local examination office. However, using Shibboleth has a couple of drawbacks from the traditional identity management perspective:

- User data is being sent directly to the service, effectively bypassing the existing IMS and established provisioning channels on the service provider side. To store information about external users in the IMS, the service must be configured as a data source similarly to the student administration or human resource management software. This often fails due to the lack of export interfaces in the services.
- User data can only be retrieved while the user is actually using the service. Changes made to the user data after service usage has finished will not be transmitted to the service provider before the user is using the service again, resulting in inconsistent data between the identity and the service provider.
- Shibboleth provides read-only access to the user profiles. Thus, on the one hand complementary technologies are required, e.g. for the transmission of exam results. On the other hand, services which need to store additional data in the user profile are forced to maintain their own local user database.

As, for example, the participants of TUM's e-learning courses shall be able to also use the related mailing lists, file space, and discussion forums, it will become necessary to connect the e-learning software as a source system to the IMS, resulting in severe administrative overhead.

3.3. User Correlation to Support Data Quality Management

As shown in the previous sections, the implementation of the LRZ's IMS as a data hub (cf. section 2) causes several challenges to the technical concept. In this section, we describe how this technical integration of various source systems causes the need for high data quality in the connected systems and how we address this issue.

As shown in Figure 2, the LRZ is importing identity data from various data sources. Since the LRZ is the IT service provider of many essential services, e.g. providing a MWN-wide groupware solution (see section 4), it must ensure that for each physical person exactly one digital identity is created. Otherwise, when having more than one account per user, the handling of IT services both for the

users and for administrators is complicated, since a user has to access to one service with login name A, whereas login to an other service would have to be done using login name B.

To assure that this requirement is met, high standards concerning data quality in the source systems are required. Many of these source systems used to store user data information only for internal use. Thus, a lot of consultations had to be done before technically connecting the various IMS to the LRZ's data hub. Figure 2 shows why this reconciliation has become necessary.

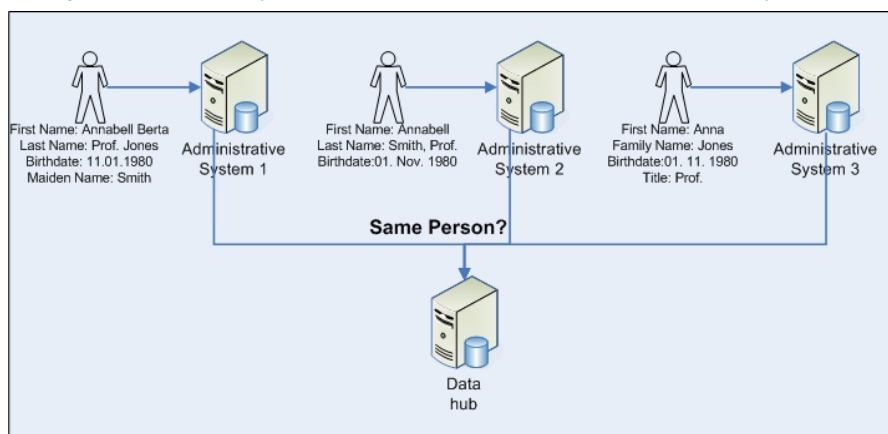


Figure 2 Data Quality Issues for Correlation

A person in this example is stored in three different source systems. Since each system was designed for a single purpose, the users' attributes were stored there disparately. In system 1 the user's full name and the maiden name is stored. In system 2 the user's data is outdated, since it does not store the actual name, and in system 3 the user's first name has been abbreviated. All systems are storing the date of birth in a different syntax and the academic title 'Prof.' is stored in system 3 as a dedicated attribute, whereas it is stored in the other systems as part of the name. This simple example shows that the diligence, syntax, and semantics of user data have to be addressed before the actual technical connection, since the matching of existing users needs to be done based on common attributes like first name, family name, or date of birth:

Diligence The awareness concerning data preciseness of the administrative staff had to be raised. For internal usage the entries' accuracy might have been appropriate, e.g. for TUM's course planning system UnivIS, where the users' exact data (like storing all first names or the date of birth) has been of second concern. Now, access to IT services is granted based on special entries as described above. Hence, the administrators need to be aware of the implication wrong entries might have for the users, if for example the booking of a course A in the e-learning system is only granted for users which have completed course B before.

Syntax The syntax of the various source systems needs to be known in advance, since a matching which is based, for example, on the date of birth would fail, if the systems are using different formats as shown in the example.

Semantics The semantics of the user attributes have to be configured. As shown in the example, the attribute "first name" could contain all first names, only a single name, or even an abbreviation. This issue is also important concerning academic titles (Prof., Dr., etc.), titles of nobility ('earl of', etc.) or name affixes ('van de', 'von der', etc.). Some systems are storing these attributes as part of the family name whereas others are storing it in dedicated fields.

To ensure that only one account per user is generated, we have implemented a correlation algorithm for matching users based on their common attributes. There we are regarding first names (all of them, ignoring symbols like '-' or ','), family names (including maiden name, since a change of one's surname, e.g. due to marriage, could have not been performed in all systems, also ignoring symbols like '-' etc.) and the date of birth, which is converted beforehand into a consistent format, of the source systems data. The correlation is done before storing the user data in our import and export directory (Figure 2). If a match is found, the user's account in our directory is enriched with the missing attributes; if no matching has been found, a new account will be generated. This implementation reduces manual overhead through manual creation of user accounts.

There inherently is a residual risk due to typing errors. Therefore, a process for manually correcting the automatically created entries is needed. Furthermore, it is not reasonable to e.g. process common typing errors (like type '4' instead of '7', or 'Smyth' instead of 'Smith', or 'Annabel' instead of 'Annabell', etc.) due to the implementation overhead. Any failing of our algorithm can thus only be caused by insufficient entries in the connected source systems. As a consequence, these source systems have to do such corrections locally on their site. These changes will be updated with the help of our algorithm the next time the user's data is sent.

To improve the quality of data in our connected source systems, we have extended this algorithm by checking for possible errors before we start the correlation. For example, we are sorting out staff younger than 14 years and are comparing (all) the names to find missing entries as shown in the above example.

Meta-directory implementations are causing significant challenges to establish unified access solutions, but they are also a means to improve data quality management. Having a single source for identity and access infrastructures, correlation of user data allows finding errors or inconsistencies in source systems.

4. CONCEPTS AND TOOLS FOR DELEGATED ADMINISTRATION

The LRZ in its role as IntegraTUM project partner is responsible for the project's sub-tasks directory services and meta-directory, storage and file systems, and email and groupware.

It has been decided to use NetApp appliances as storage systems for home directory, project, and mailbox stores. Microsoft Exchange has been chosen as the strategic TUM-wide email and groupware solution. In consequence both services (storage and email) are based on Microsoft Active Directory Services (ADS). Therefore the LRZ as the IT service provider for all HEIs within the Munich scientific network designed and implemented a hierarchically structured and multi-subscriber capable ADS Directory Information Tree (DIT) to prepare the offering of these services for the complete MWN. Regarding the improved auditing abilities and better collaboration possibilities, a high integrated and high available MWN-wide single domain architecture has been chosen. This architecture has been preferred in favor of a loosely coupled multi domain forest interconnection. The chosen architecture offers extended flexibility regarding inter-HEI (MWN-wide) collaboration, e.g. between LMU and TUM in medical sciences or physicians' cooperation projects together with Max Planck Gesellschaft Garching (MPG), for example within the universe cluster of excellence (see Universe Cluster website (2008)).

Within the single domain DIT each HEI is represented by an organizationalUnit (ou) LDAP entry. Beneath the institutions' ou sub tree, each organization's structure is represented by two consecutive ou containers fulfilling the HEI internal administrative needs. The maximum predefined intra-HEI layer depth is two. For the TUM the first level represents the faculties (e.g. the ou "Informatics" or "Physics") the second level the chairs, which are represented by the ou "Chair1" in Figure 3. Regarding the IntegraTUM project's scope the administration processes within the TUM are focused in this section.

The administration processes at the TUM have grown over years. Nearly each organizational unit has its own IT services and administrators. IT system administration in the past has often been done by researchers whose core competence obviously is not the management of IT systems. This fact has often lead to suboptimal managed IT systems resulting in decreased availability on one hand as well as additional work and overtime for affected researchers on the other hand.

4.1. Background

One of the goals of the IntegraTUM project is the optimization of IT system management processes regarding the availability, the distributed administrative efforts, and the resulting service quality.

The idea is to recentralize the IT systems' operation (hardware and administrative work regarding the IT services themselves), and to offer in parallel each unit or department access to manage their specific data and resources on their own. This results in various advantages: For example, researchers are released from time consuming and off-topic tasks; IT experts are responsible for the central user management and for running and managing the core services. The core services include globally shared resources like network equipment, storage systems, or mail systems regarding tasks

like backup, upgrade, security, and patch management, as well as the identity management tasks (e.g., import processes). In the past all these tasks had to be handled by each department or even on basis of individual chairs. Work has been done twice or even more often, which resulted in differing and sometimes suboptimal data quality. MWN-wide cooperation and collaboration, e.g. using groupware solutions that offer shared mailboxes and calendars, was hard to realize. Even any kind of TUM-wide (between organizational units) collaboration has been a challenge. User account information of unit X not necessarily has been present in unit Y. For co-operations, this often has led to naming collisions caused by non-existing or just incompatible namespace definitions, which in the past resulted in at least doubled user account objects with different names (login names) and passwords, depending on the local security policy.

The local administration tasks have changed in a way that local administrators

- benefit from using the centrally available unique user account information.
- do no longer need to know all the details about IT systems and their maintenance.

In consequence, the local administrative workload has been reduced. For most of the services, administrative work has been replaced by a single, simple task: the assignment of already existing, locally accessible user objects (locally known users) to globally available and running resources.

However, not all IT resources (e.g., locally demanded and highly specialized systems) have been recentralized. While these resources have to be maintained locally, the existing user account information can be accessed nowadays. Depending on the platform technology (e.g. Microsoft's Active Directory Services), TUM's IT administration concept also includes these decentralized resources in the existing infrastructure and administration processes. Local administrators of local resources can take advantage of the identity management infrastructure. They benefit from the unique user names when granting permissions on their resources. The end users profit from the seamless integration and the improved usability (e.g. unified login and single sign-on).

4.2. Management of Centralized Resources

Any resource, its description and name are well defined, globally known and globally made available. The central resources are internally "divided" into separate parts (DIT branches), according to the organizational units' structure. Each of these separated areas is accessible from remote. The access depends on the roles a user belongs to within an organizational unit. While administrative role owners are allowed to manage access permissions to an organization's resources, user role owners are just allowed to use the assigned resources according to the granted permissions.

4.3. Decentralized Administration

Although the assignment of single user objects to available resources is supported, it is not advised because it would be rather ineffective, error-prone and not reusable at all. A better and more elegant solution is a two step process. In the first step, the local administrator creates a group and assigns access permissions for a resource to this (initially empty) group object. In the second step the administrator assigns user objects to the group objects. This process offers various advantages. The permission set is reusable; adopting permissions can be done quickly for a huge amount of users. The same group can be used to grant access to different resources (e.g. project file space and an according mailing list or calendar access) without repeatedly assigning any single user to this second resource. Adding a new user to a group will grant access to n assigned resources. As the group object itself is stored and managed on a central system, it is available to all connected systems (even local resources of other organizational units).

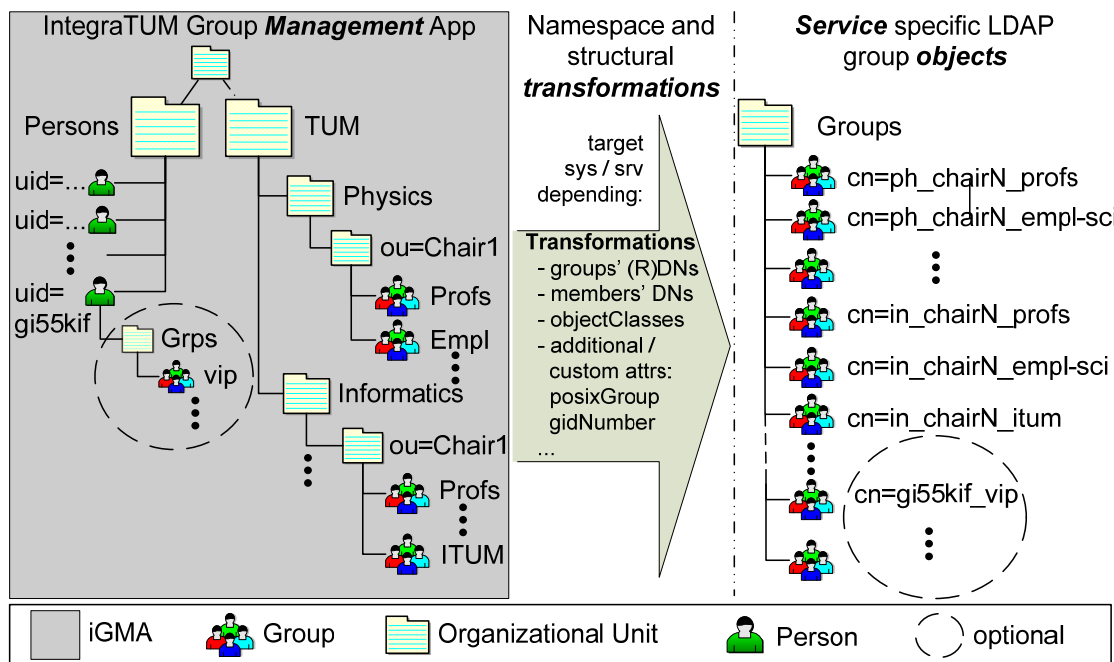


Figure 3 iGMA - IntegraTUM's group management application backend presented in (Pluta, D. (2007)) offers hierarchical management access to group objects (on the left side) and provides consumer side specific transformed group objects for connected IT systems (on the right side).

Independent from the locally created ones, managed and globally made available groups, as well as predefined and/or automatically created groups offer further advantages. Depending on the user data sources' data quality, these groups can at least be created in advance and possibly even populated automatically using filters on user object attributes (e.g. employees, employees of unit X, students, students of faculty Y in semester N). Regardless whether these predefined group objects are initially empty or populated, they get assigned default access rights. The resulting advantages open out in minimized local administrative overhead. Most common group objects are already existing and usable. They do not have to be created first. Because they are created automatically, their names are predictable. By using a name concept that offers a transparent rule set to infer global group names, it is possible to pre-assign permissions to individuals of kind X without the need to locally administer any group at all.

Establishing a multi-subscriber capable IT service is a technically complex task. Depending on the target platform many possibly in the future occurring side effects have to be taken into account in advance (e.g. adding or renaming organizational units or any kind of resources). After the successful design of a technical concept the settling of all organizational and political doubts is at least of comparable complexity.

4.4. Tools

In the following we present a short overview about the most important tools we are currently using within our IT infrastructure. Commercial tools like Microsoft Active Directory or Novell's eDirectory as well as open source based and custom developed software is currently deployed. The deployed services are under continuous development. Regarding the user support and change management we have established a service desk to channelize the process and communication flow.

Open Ticket Response System Within OTRS each participating organizational unit is represented by a queue. Each queue is responsibly managed by service agents. Currently more than thirty queues that are accessed by overall about 90 service agents have been setup. In (Hommel, W., Knittl, S. (2007)) the details regarding the design, the supported processes (e.g. incident management), and technical details are discussed.

eDir Directory Service Our complete identity management core infrastructure and implemented core processes are based on Novell's eDirectory services and in addition also Novell's Identity Manager software component (formerly known as DirXML). The combination of these two tools offer event based directory operations to support just in time provisioning (including credential information) of all connected systems.

MWN-wide Active Directory For the Microsoft Windows platform we have chosen the single domain active directory (AD) model. Internally the organizational units are represented as organizational unit LDAP entries (*ou*). In combination with the previously described name space definition the single domain model offers various advantages regarding more efficient administration processes, support for inter domain delegation, security audits, and last but not least automatic conflict checks that avoid ambiguity regarding a resource's name. Based on the AD integration of the MWN-wide NetApp storage system and Microsoft's Exchange messaging and groupware solution all AD internal users are able collaborate by sharing resources like project shares and calendars for example.

IntegraTUM Webdisk This tool is a free web application that gives direct access to a file server. It is written using Java Servlets and the jCIFS library. Supported file servers are Samba, MS Windows and NetApp OnTAP and those based on the CIFS-protocol. The IntegraTUM WebDisk offers a flexible way to access files remotely without the necessity to use an active directory integrated client system (e.g. internet browser software offered in an internet café). The application is published as open source software and can be downloaded from the IntegraTUM WebDisk website (2007).

Group-Management Backend In addition we have implemented a powerful LDAP backend based on OpenLDAP that supports delegated group administration (flexible group membership assignment, see (Pluta, D. (2007))). The frontend component for handling the backend-internal, hierarchically stored and managed group objects is integrated into our modular web based management framework prototype (group management module). As outlined in Figure 3 the management tree shown on the left side illustrates the management view regarding the group objects. On the consumer side (e.g., storage and groupware systems) the group objects are represented by and accessed as a flat data structure, illustrated on the right half of this figure. The naming concept rules are applied during the transformation process in-between.

Guest-Management Middleware and Frontend User objects are created during the import of employee data sets from SAP HR and student data sets from the HIS-SOS system. Users not listed in either of these data sources, such as the university's guests, have to be given selective access to IT services, too. Thus we implemented a middleware based on XML-RPC services that is able to create guests user objects according the requirements of the identity management system. The middleware functions are called by a web-based GUI implemented using Microsoft Office Sharepoint Server technology. Due to less dependencies the service oriented architecture simplifies the GUI development and offers the possibility to independently replace the frontend, backend, and user interface by another technology. Currently it is planed to integrate the guest management frontend into the new campus management system called TUMonline.

Web-based Management Framework Prototype We have developed an extensible and modular web framework prototype based on Java and JavaServer Faces technology. The frontend is designed as single point of contact offering a common look and feel for all kind of IT service related tasks. Management operations that belong together (e.g. all tasks regarding the group management: search, add or modify) are bundled in separate modules. Internally the framework assigns one or more unique role definitions (e.g. the roles group-mgmt-admin, group-mgmt-readonly...) to each module. For an authenticated user a module or even some modules' sub-tasks can only be accessed if she has been assigned as role occupant of the same role(s) the module has been assigned before. The amount of role objects is not limited so the framework is able to be populated and extended by independently developed modules with individual authorization. The role objects themselves are stored within the above mentioned group-management backend and they are manageable using the web management framework's group management module itself.



Figure 4 The IntegraTUM AdminGUI (IAG) is a modular web-based management frontend implemented using JavaServer Faces (JSF)

The different modules are accessible depending on the current user's role ownership. Standard access is given to the "Self-service" module, so at least each user is allowed to change her password. Figure 4 shows an internal view, behind the login dialog. As can be seen in the navigation bar, a user possessing six roles is currently logged in: this user is allowed to access the modules "Administration", "Self-service", etc. As can be seen in Figure 4, the module "group management" which itself offers four module specific sub tasks has been selected. For example the "group management module" allows a local administrator to manage group objects stored within a LDAP backend.

The frontend is completely localized. It currently offers English and German language support. However, it is easily extendable regarding additional language support. Due to the separation between program code (internal logic) and translation data no programming skills are required to add further translations. As another important requirement the front end's layout and design is based completely on cascading style sheet (CSS) technology. Thus the layout can be exchanged and customized the most flexible way (e.g. adapting the design regarding cooperate identity guidelines). There is no need to access the internal program logic to change the web GUI's presentation layer. The continuous use of CSS offers further advantages; imagine optimized content display and application handling to offer accessibility. Different media types (e.g. standard browsers, browsers offered on mobiles, and PDAs or text only clients) are currently supported.

5. CONCLUSION

Within the higher education institutions in the greater Munich area, over 100,000 user accounts have to be managed. IT services offering access to e.g. digital libraries or learning management systems as foundation for cross-institutional research or collocated study courses need a profound technical and organizational solution. In this paper, we presented our technical solution, which is based on the Identity Management Systems operated at LRZ on behalf of MWN member organizations including TUM and LMU. Based on the sketched IT system components' interaction we discussed the improved support to automate several of our organizational core processes, which results in more reliable data and enhanced process quality.

In the near future, we plan to couple these systems even tighter to better support users of collocated study courses like bio-informatics. Based on LDAP technology a MWN-wide usable and standardized groupware solution has been implemented. One important prerequisite for its success

is the MWN-wide naming concept. On the basis of standardized naming schemes the applications, management processes, and tools for managing and propagating these definitions can be developed, introduced, and deployed more efficiently. Based on directory service technologies, we have also developed a concept for the decentralized administration of central resources. Hereby core IT services like server administration are operated by IT specialists at the LRZ, while allowing at the same time fine-grained administration from within the organizational units of the university.

To grant external users remote access to our IT services, e.g. in the case of international Grid projects, we have adapted our concepts and processes, for example regarding the DEISA project and the DFN-AAI Shibboleth federation. Being based on Shibboleth respectively LDAP, both solutions allow external users access to the required resources, helping to fulfill the goal of the Bologna process of supporting the mobility of students and researchers. Our further work here concentrates on extensions to the DFN-AAI schema in order to support the requirements of learning management systems. In the Grid area, we are taking part in the IVOM project to better combine the infrastructures of D-Grid and DFN-AAI in future.

The technical approach introduced in this paper would not have been realizable in a high-quality manner without optimizing some of our organizational processes. Thus, we are currently extending our IT Service Management processes, even across organizational borders. This helps us to better fulfill our users' needs regarding the access to our IT infrastructure. In the future, we intend to continuously strengthen our system integration to establish further possibilities in concern of cross domain collaboration towards the goal of Bologna.

Acknowledgments

This work was partly funded by the German Research Foundation (DFG) under contract WGI 554 975.

The authors thank the members of the IntegraTUM project team for fruitful discussions and constant encouragement. IntegraTUM is headed by the vice president and CIO of TUM, Prof. Dr. Arndt Bode (see <http://portal.mytum.de/iuk/cio/>).

The authors also thank the members of the Munich Network Management (MNM) Team for helpful discussions and valuable comments on previous versions of this paper. The MNM-Team, directed by Prof. Dr. Heinz-Gerd Hegering, is a group of researchers of the University of Munich, the Munich University of Technology, the University of the Federal Armed Forces Munich, and the Leibniz Supercomputing Centre of the Bavarian Academy of Sciences. The team's web-server is located at <http://www.mnm-team.org/>.

6. REFERENCES

DFN-AAI website (2008). *Deutsches Forschungsnetz - Authentifizierungs- und Autorisierungs-Infrastruktur*. Retrieved May 12, 2008 from: <https://www.aai.dfn.de/>

Universe Cluster website (2008). *The Cluster of Excellence for Fundamental Physics*. Retrieved May 12, 2008, from: <http://www.universe-cluster.de/>

Hommel, W., Knittl, S. (2007). *SERVUS@TUM: User-centric IT Service Support and Privacy Management*. In proceedings: 13th International Conference of European University Information Systems. EUNIS 2007, Grenoble, France, June 2007.

IntegraTUM WebDisk website (2007). *IntegraTUM WebDisk*. Retrieved May 12, 2008, from: <http://sourceforge.net/projects/webdisk/>

Pluta, D. (2007). *ACL Design Behind IntegraTUM's Decentralized and Delegable Group Management*. In proceedings: First International Conference on LDAP - LDAPcon 2007. UpTimes No 3, Cologne, Germany, Sep. 2007, pp. 27-35.