SERVUS@TUM:
User-centric IT Service Support and Privacy Management

Silvia Knittl[1] and Wolfgang Hommel[2]

[1] Technische Universität München, Dept. of Computer Science, Boltzmannstr. 3,
85748 Garching, Germany, knittl@tum.de
[2] Leibniz Supercomputing Center, Identity Management Team, Boltzmannstr. 1,
85748 Garching, Germany, hommel@lrz.de

Abstract:
The Technische Universität München (TUM) has been awarded as one of three German elite universities. To better support students, researchers, and guests, as well as placing them in control of the use of their personal data in a growing number of inter-organizational projects, we designed *user-centric* solutions with a strong focus on both the organizational and technical aspects. In this article, we first present the concepts and current realization status of our university business process driven recentralization of the IT service support; as many essential parts of TUM's IT infrastructure are handled by the Leibniz Supercomputing Center, which is the common computing center of the higher education institutions in the Munich area, emphasis has been put on cross-organizational processes and their tool support. Then, we present a user-centric privacy management tool that has been implemented for the Shibboleth middleware and enhances the users' control of their personally identifiable information.

1. Introduction

Efficient research and education obviously require a modern IT service infrastructure. At EUNIS 2006, we reported about the project IntegraTUM, which comprises our ongoing work to recentralize the key IT services at Technische Universität München (TUM) in order to **provide them more cost-efficient and facilitate the cross-institutional collaboration**, which has suffered from a larger number of decentralized established services such as local file and email servers in the past [1]. While our central **Identity & Access Management solution has proven to be a key enabling technology**, recentralizing IT services inherently leads to major organizational challenges. In this article, we describe our user-centric approach to IT service support and privacy management, as one of our primary goals is the increased surplus from our customers' (i.e. the university staff's, students' and guests') points of view.

Our new IT service management (ITSM) infrastructure is currently being planned and realized following the reference processes of the IT Infrastructure Library (ITIL) [3]. However, as essential parts of TUM's IT infrastructure are being built and operated by the Leibniz Supercomputing Center (LRZ), special emphasis has to be put on cross-organizational processes, which are not covered by best practices collections such as ITIL yet, and are subjects of our ongoing research in the ITSM field. Furthermore, privacy and data protection issues, which often have been neglected at small decentralized IT service installations, become a key issue for EU privacy law conformity, compliance requirements, and user acceptance. Based on the authentication and authorization middleware Shibboleth, which is currently being rolled out as part of the German NREN's (National Research and Education Network) authentication and authorization infrastructure DFN-AAI [4], we have designed and prototypically implemented an enhanced component that allows a very fine grained control of the flow of personally identifiable information (PII) by the DFN (German Research Foundation), the TUM, and especially by each user.

This article is organized as follows: In section 2, we summarize our previously decentralized IT service support and its deficiencies, followed by a detailed explanation of our enhancements and the overall new architecture in section 3. Our primary motivation for improving the user centric privacy management is outlined in section 4, along with a summary of our concept for Shibboleth and a presentation of our implementation. Finally, section 5 summarizes this article and gives an outlook to our future work.

2. Why decentralized IT service support does not scale, but instead confuses the users

Previously, **TUM's service support infrastructure has grown decentralized** along with the services. Figure 1 shows an incomplete selection of different IT services according to TUM's organizational division. IT services are provided by the departments, internal administration, through projects or external suppliers like LRZ. Those IT services vary from e.g. labs at the physics department to storage services at the LRZ or maintenance support for the staff's PC from the administration.

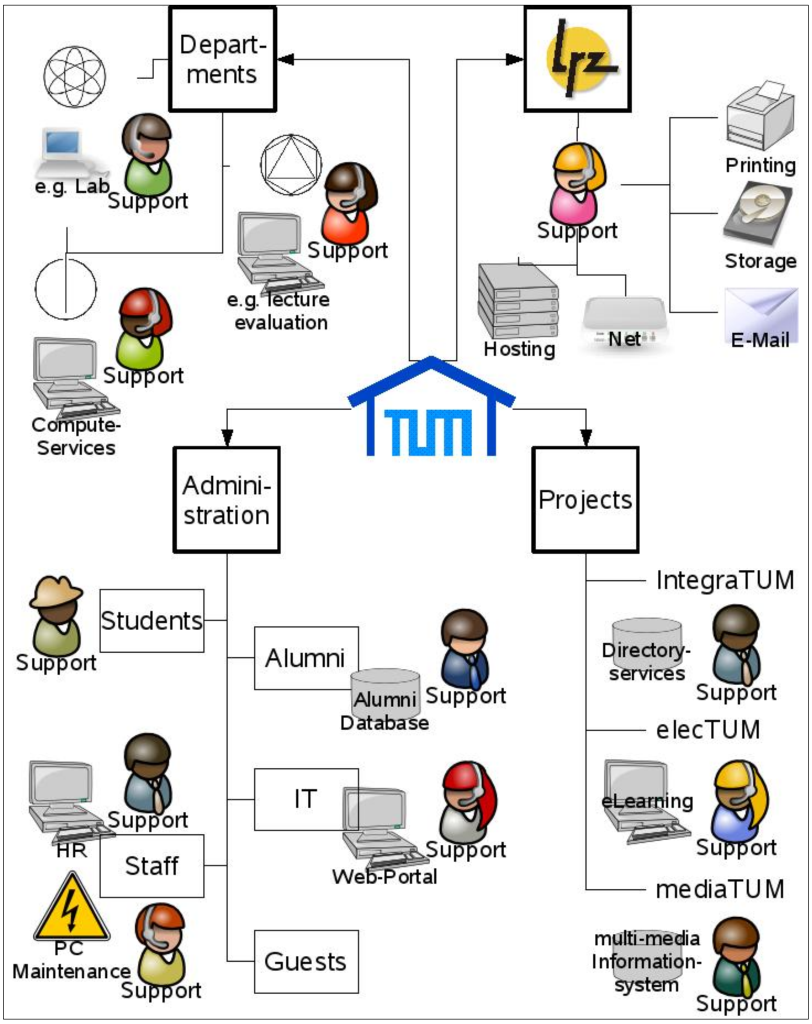

Figure 1 TUM's previously decentralized support infrastructure: which support would you choose?

As shown in the top left-hand corner at Figure 1, chair computer labs for student practicals were built and operated independently of the faculty's student computer pool. Students were assigned a new user-name and password for al-most each of their curricular activities, leading to a lot of login failures and forgotten passwords, which in turn resulted in additional cumbersome work for the service administrators.

Often, chairs did not have dedicated positions for administrators, and instead young researchers had to stand in; as many of them would rather have focused on their teaching and research work, the **quality of service often was suboptimal**: Especially concerning system and network security, many incidents which caused complaints had to be accounted for and required additional labour.

Even for many employees it often was not obvious whether an IT service was being provided by their local chair, their faculty, one of the university's facilities, or the LRZ. For getting support, people had to find and choose from more than 15 different email addresses and the same quantity of phone numbers. Consequently, incident reports and questions were only seldom directed to the appropriate contact person in the first step and instead had to be forwarded until they eventually reached someone who could help.

A test conducted in the early days of the project showed that it took more than a week and more than seven email forwards until the question returned to the original sender without anyone being able to solve the issue. **An investigation of these symptoms** as part of a diploma thesis [15] revealed the following nuisances:

- Services were often provided by individuals who had either not enough time for supporting the services due to their other work load, or were only technically oriented and not really interested in supporting their users.
- The **need for support was often recognized too late** when building a new service, so in many projects there was not enough time or money left to set up an adequate support infrastructure when the service had already gone into the production phase.
- There was no central authority or registry concerning the introduction of new services; thus, many of those who tried to help the users did not even know a service existed until someone reported a problem with it, causing additional overhead to figure out who was responsible and could help. This took even more time whenever someone's assumption that someone else could help turned out to be wrong.
- Except for the few central services, for example in the university library and the LRZ, existing support staff had not been trained and often was unable to answer user questions in an adequate, not too technical manner.
- **No support processes had been established**, especially concerning the cooperative work between two or more service providers to solve a common problem.
- **Hardly any of the support instances was appropriately tool-supported**. For example, only the LRZ and very few of the faculties were using a trouble ticket system and a known error database. This not only resulted in knowledge getting lost over time, but also led to sometimes multiple and different answers to the same questions by different people who did not know that someone else already took care of the problem. Furthermore, the lack of standardized incident reports often led to missing information which would have been necessary to help; for example, only based on the sender's email address, it was impossible to decide whether the user was a student or employee, especially when people didn't use their university email addresses, but their private ones.

Obviously, the overall service support quality level was quite low. Thus, when the IntegraTUM project [2] started to recentralize services such as email and file servers, it quickly became evident that the project's success largely depended on the establishment of an efficient and comprehensive service support infrastructure.

3. SERVUS@TUM: University IT Service Management based on ITIL concepts

SERVUS@TUM (in Latin, servus means "a servant"; the word "service" is derived from it. In Bavaria "Servus" is also a friendly greeting like "Hello") is a sub-project within the IntegraTUM project. The goal of SERVUS@TUM is to introduce a **comprehensive service**

support infrastructure; its main focus is the field of establishing an Incident Management Process and a central Service Desk (SD) – based on the best-practice framework ITIL [3]. ITIL is a composition of books as guidance for holistic ITSM. It was originally developed by the Office of Governance Commerce (OGC) on behalf of the British government, but is now further developed with various worldwide partners within the project "ITIL Refresh" to ITIL Version 3 (to be published by end of May 2007). The orientation of ITSM processes according to ITIL lays the foundation for a certification with the new standard ISO 20000 [12]. Furthermore, the processes of ITIL do support also the objectives of common IT governance frameworks like COBIT (Control Objectives for Information and related Technology) [13]. Therefore, ITIL - when used in conjunction with an IT governance framework - provides a hierarchy of guidance for IT Management. The following advantages can be achieved by the introduction of ITIL:

- **improved quality of service through usage of proven, best-practice processes**
- improved customer satisfaction through a professionalized service delivery
- improved productivity
- improved assignment of qualified staff
- integrated and centralized processes
- verifiable performance indicators

The aim of the Incident Management Process is to restore service operation in the case of an incident as quickly as possible. An incident is defined as any event that (potentially) provokes an interruption or a reduction in the quality of service within the agreed standard operation of a service. This process ensures that breakdowns of services are as short as possible. To achieve this, the actions taken by the Incident Management Process are registration, detection, classification and diagnosis of incidents. Furthermore, the SD supports Incident Management by providing the required inputs. The SD guarantees the reachability of the IT organization by providing an institutional interface for the users, effectively being a Single Point of Contact (SPoC). The members of SD accept the users' incidents, register them, often immediately resolve them based on a known-error database or pass them to the appropriate specialists if necessary.

To support all these tasks, the introduction of a trouble ticket system is recommended [3]. Requirements of such a system are among other things:

- The possibility of automated logging of incidents. This ensures that none of the incidents are getting lost, like it could have been possible in the traditional way, where incidents were written on memos sometimes.
- The passing of incidents to other units needs to be flexible, because the various co-workers of the IT at TUM are located in different geographic places and organizational units.

The scope of SERVUS@TUM is to introduce a multi-tier support infrastructure for TUM. Figure 2 shows our organization with the central SD at 1st level and the supporting units at 2nd level. The main focus of this SD lays within IT technology. The coordinated support of students regarding issues of their studies lays within the scope of the Studenten-Service-Zentrum (SSZ). To realize the above described tasks of the 1st level support – starting with registration of incidents - we have already provided an email interface (it-support@tum.de) for submitting incidents to the central SD. Incoming incidents are already predominantly solved by the 1st level, or they are solved in a coordinated manner with 2nd level support units. Doing this, we have already achieved first improvements in the duration of incident handling since the users do not have to care to find the appropriate support units – as described above – but the co-workers of the SD.

Our next step to endorse the SERVUS@TUM's technical infrastructure is the introduction of the trouble tickets system Open Ticket Request System (OTRS) [14]. OTRS is an open source system, which has been proven as an adequate means for our purpose in an earlier evaluation. OTRS allows Incidents to be issued in a channelled way by either email, web or phone. This system has been already in use at a few TUM organizational units. At present we are testing OTRS within SERVUS@TUM with the 2nd level support units of the physics department, with student's administration and with the LRZ as an external IT service provider. All the knowledge that we gain within this project phase will be integrated in our concept, and the system will be adapted to be able to cover all requests to integrate all other existing support units.

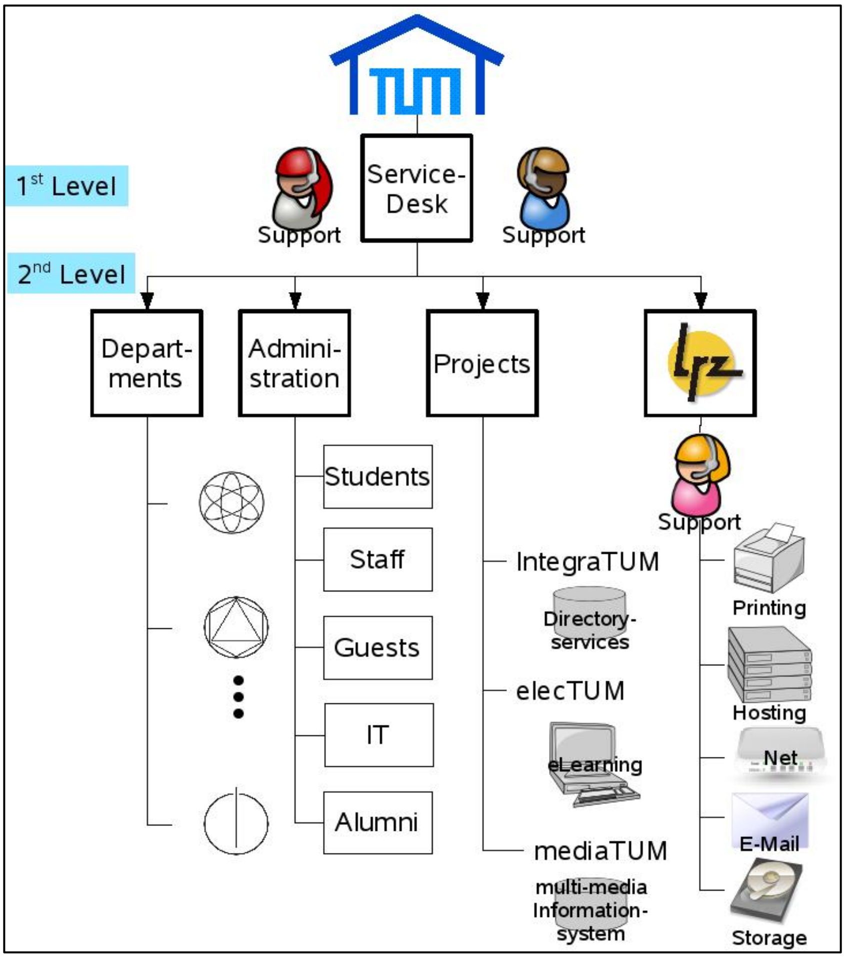

Figure 2 TUM Organization and IT-Services with central support

Even in this short period of testing **we were able to achieve the general advantages referred in ITIL literature**. By introducing the SD we have created a well defined contact interface for all sorts of technical incidents. Thus, all incidents are now recorded within OTRS. Furthermore, trivial issues like login problems are prevalently solved within 1st level support, thus unburdening specialists in 2nd level and therefore improving their working situation. Moreover the overall duration of incident solution is decreasing since not users have to go from pillar to post any longer to find the appropriate supporting unit but the co-workers of SD are taking care of incident routing.

The orientation of the concept of TUM's SD on ITIL's best-practises allows us a faster implementation. However, **ITIL does not provide any solutions in case of cross-organizational processes** like TUM's identity management, where various TUM internal organizational units as well as external suppliers like LRZ are working together in providing the service. This is why we are continuing **our research in the field of federated ITSM**. Especially in the domain of Configuration Management, which is according to ITIL, responsible for providing actual data about the configuration items in use with the help of a

Configuration Management Database (CMDB). For cross-organizational processes **the implementation of a federated Meta-CMDB will definitely be necessary**. Until now, there are no approaches for this issue within ITIL.

According to our experiences with the TUM projects IntegraTUM, mediaTUM [17] and elecTum [16], further positive synergies can be gained for users of our IT infrastructure if we continue on the topics integration, collaboration and co-sourcing. Within the scope of IntegraTUM, customized identity information is provided for different applications. The goal of mediaTUM is the development of a multi-media information- and archive system and elecTUM is providing TUM's central learning management platform. To further support these projects but also the users of this support infrastructure the following challenges are highly relevant for SERVUS@TUM:

- **Integration**: By integrating new technologies like instant messaging or voice over IP, the SD will be established as a multi-medial base for communication and information. Information collected here can help IT management for decisions in further improving the existent services or introducing new ones.
- **Collaboration**: These newly integrated technical facilities allow better collaboration of the various co-workers of SD within the various organizational units respectively locations as well as the collaboration amongst the users. The planned administration of arbitrary groups within IntegraTUM allows the establishment of mailing groups for mailing lists or learning groups within the learning management platform provided by elecTUM.
- **Co-sourcing and shared services**: The experience within the project IntegraTUM has proven that the partnership with external service providers like LRZ is effective. By extending this co-sourcing relationship more advantages will be achieved. The provisioning of IT services as a shared service amongst TUM internal organizational units is improving service quality as a result of reduced failure rates and also is improving efficiency. With the help of an authentication service, provided by IntegraTUM, we achieved a raise in efficiency since the local user administration became spare as well as an increase in quality of data, since changes in the provided data are now allowed only using dedicated applications. The upcoming **ITIL version 3 will provide us with more design guidance** for co-sourcing as well as the implementation of shared services.

For the further development of the SSZ, we plan to transfer the know-how gained with the conception and implementation of the SD to the technical infrastructure of the SSZ. Thus the user-centric approach can be applied to the TUM as a whole. The scenarios of distance learning, introduced in the next section, confirm that the concepts developed within SERVUS@TUM can be used as a base for a more global approach. The cross-organizational aspect of ITSM becomes also evident in the following section, where solutions for privacy management within Federated Identity Management are being described.

4. User-centric Privacy Management: Using Shibboleth in complex Identity Federations

DFN is currently establishing its **authentication and authorization infrastructure**, DFN-AAI [4]. Based on the Shibboleth middleware, which already has been deployed in several other European countries, its primary use cases will be the cross-organizational use of library resource access, distribution of licensed software, and e-learning.

The use of Shibboleth in combination with learning management software is also planned for several study courses which are shared between the TUM and several other universities in Bavaria. Several hundred students enrolled in those study courses are members of multiple universities and partially require access to TUM's IT services without being registered in TUM's identity management solution. By Shibboleth-enabling certain services like TUM's web portal and the learning management system, these students could use those services transparently without having to redefine the underlying business processes from scratch (cf. figure 3).

However, using learning management systems via Shibboleth requires the transfer of personally identifiable information such as date of birth, matriculation number and sometimes even examination results, which exceeds the amount of data typically, transferred using AAIs. Thus, this scenario is **more complex regarding privacy and data protection** than most previous Shibboleth-based infrastructu-res.

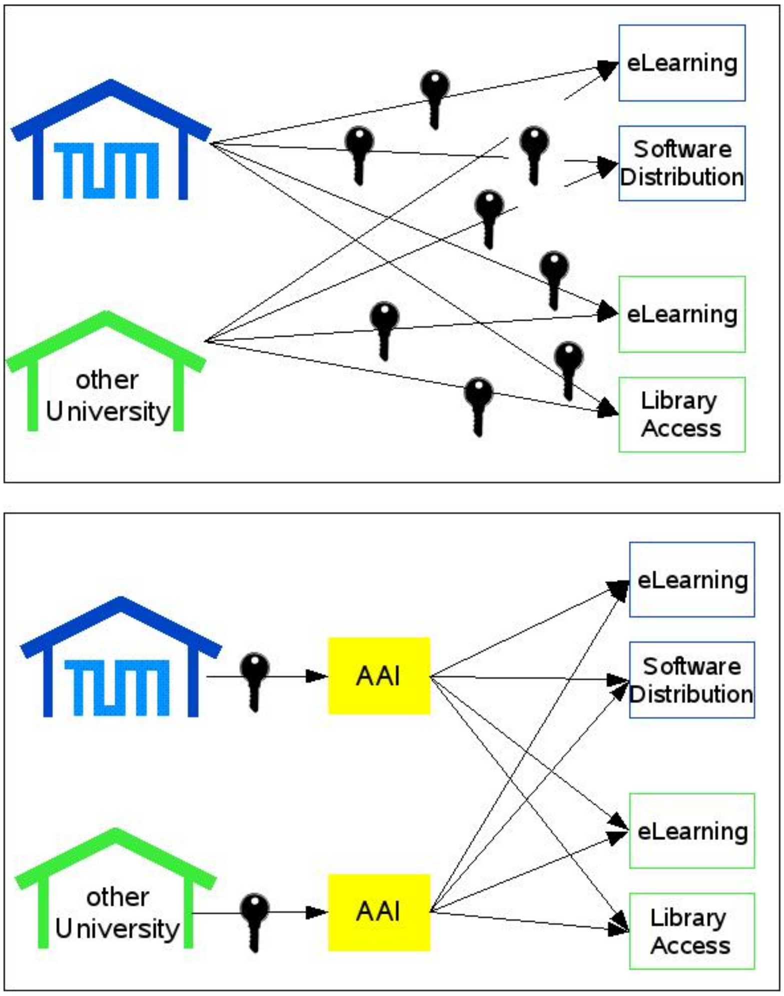

Figure 3 Without Authentication and Authorisation Infrastructure (AAI) (above) and with AAI

When using Shibboleth, the so-called Service Providers can request user attributes from the so-called Identity Providers (IDPs), which are the students' home universities. A rule-set called Attribute Release Policies (ARPs) controls which user attributes may be sent to which Service Providers. There is one Site-ARP, in which the IDP-wide defaults are specified, and optionally one ARP per user, in which these defaults can be restricted even further for individual users.

Shibboleth ARPs must be formulated in a rather simple XML dialect; for users, a few dedicated web-based graphical management frontends exist, e.g. the popular Shibboleth Attribute Release Policy Editor SHARPE [5]. However, Shibboleth's ARPs have been designed and are sufficient for rather simple scenarios in which often only a person's entitlements are transferred to the Service Providers. In the e-learning scenario, Shibboleth's implementation of ARPs exhibits certain drawbacks:

- There can be only one site-wide and one user-specific ARP; however, ARPs cannot be specified federation-wide or group- or role-specific.

- Attributes and Service Providers cannot be grouped within ARPs. Thus, when four attributes should be transferred to five distinct Service Providers, 20 separate rules must be specified and maintained, leading to **significant administrative overhead** in complex scenarios.
- Shibboleth does distinguish between services, but not between different use cases of the same service. Thus, all attributes specified in ARPs are sent to the Service Provider independent of whether the user is actually using the service or just browsing the service's web site for information; considering the data being transferred in the described scenario, this can be a significant privacy breach.
- Finally, there is **no support for obligations**, such as the demand to delete a user's personal data e.g. at the end of the current semester.

In [6], we have presented a concept that demonstrates how the policy language XACML (eXtensible Access Control Markup Language, [7]) can be efficiently used to model and enforce Attribute Release Policies that comply with the requirements of the state-of-the-art in Federated Privacy Management. Meanwhile, this concept has been implemented as part of a diploma thesis [8]; subsequently, we summarize the results of this work:

- An URI-based namespace concept has been implemented which supports
  - Globally unique user attributes and multiple identities, roles, and groups per user.
  - Service providers which can offer multiple services and have multiple sub services, use cases, and attribute requesting purposes per use case.
  - The grouping of attributes and services in order to reduce the total number of policies that have to be specified.
- Attribute Release Policies can now contain complex conditions, for example making use of contextual and environmental data such as the current date and time or the student's current enrolment status.
- Obligations, such as sending an email to a user whose attributes have been requested or writing to individual log files which can be viewed by the users, are fully supported. For long-term obligations, such as deleting the information e.g. 90 days after the last service usage, an interface to privacy management systems such as EPAL [9] has been designed.
- Attribute Release Policies can be arbitrarily prioritized, so for example the user's ARP always overrides the site-wide ARP or vice versa, depending on the scenario's requirements. Also, federation-wide, group-, and role-specified ARPs are supported. The site-wide ARPs can also be administered in a decentralized manner, so for example policies defined for the whole IDP can be refined by department and working group specific ARPs.

In the implementation, Shibboleth ARP's engine has been extended by a XACML policy enforcement point, which makes use of Sun's reference implementation of a XACML policy decision point [10]; the ARPs are stored in an OpenLDAP-based directory service, which uses LDAP Access Control Lists to control read and write access to the policies. Optionally, XML signatures can be used to ensure the integrity of each ARP.

The open source implementation is available as patch for Shibboleth's Java source code (version 1.3c) from [11]; future work will focus on the implementation of a suitable graphical XACML ARP editor for the users: As XACML is much too complex and error-prone to be used directly by the customers, its complexity must be hidden by means of an adequate management frontend. Existing Shibboleth ARPs can be converted lossless to XACML ARPs by using an XSLT stylesheet.

5. Conclusion and outlook

In this paper, we presented TUM's user-centric approaches towards IT service support and privacy management. Based on the ITIL best-practice concepts, several adaptations had to be made to **support cross-organizational processes**, especially concerning the service desk, the Incident Management, and the Configuration Management. We introduced our new **layered support architecture** implemented within the project SERVUS@TUM. Formerly highly decentralised and self-contained support workflows are now integrated within a single and **comprehensive incident management process** resulting - despite being in project state - in a reduction of incident handling time, IT professional's relief from routine jobs and increased customer satisfaction. Within SERVUS@TUM we further plan to integrate new techniques to enable TUM's Service Desk to become a **central information and communication platform** to support IT management and also enable collaboration amongst co-workers and users of our IT infrastructure like e-learning platform.

The use of Shibboleth for rather complex e-learning scenarios has led to the need for more sophisticated and easier to administrate Attribute Release Policies; based on the standardized policy language XACML, we have implemented a prototype that supports multiple roles per user, arbitrary complex conditions and obligations. It will be tested and improved in TUM's identity provider for the German DFN-AAI federation.

Our future work will focus on a more complete realization of ITIL-based processes at the TUM and the LRZ. Again, the **cross-organizational aspects pose challenges for which no suitable best practice solutions are known yet**. Here we plan further research in the area of **federated Configuration Management** since it provides an information repository regarding the whole IT infrastructure and thus supporting all ITSM processes. Our work on Shibboleth will continue as the DFN-AAI federation is going into production phase; concerning the XACML-based Attribute Release Policies, work has especially be put into the design and implementation of an intuitively usable graphical editor. Also, **users will have to be sensitized** and made aware of their privacy-related options, along with the specification of suitable default policies for the popular services.

Acknowledgment

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

The authors thank the members of the IntegraTUM project team and the Munich Network Management Team for valuable comments on previous versions of this paper. IntegraTUM is headed by the vice president and CIO of TUM, Prof. Dr. Arndt Bode (see http://portal.mytum.de/iuk/cio/).

The MNM-Team, directed by Prof. Dr. Heinz-Gerd Hegering, is a group of researchers from the University of Munich, the Technische Universität München, the German Federal Armed Forces University in Munich, and the Leibniz Supercomputing Center of the Bavarian Academy of Sciences. The web server of the MNM Team is located at http://www.mnm-team.org/.

References

[1] L. Boursas, W. Hommel: *Efficient Technical and Organizational Measures for Privacy-aware Campus Identity Management and Service Integration*, EUNIS 2006, Tartu, Estonia

[2] R. Borgeest, A. Bör: *Die IuK Strategie der Technischen Universität München – Auf dem Weg zur Digitalen Hochschule In: Informationsinfrastrukturen im Wandel. Informationsmanagement an deutschen Universitaeten*, Verlag Bock & Herchen, 2007

[3] OGC, Ed., *Service Support*, ser. IT Infrastructure Library, The Stationary Office, 2000.

[4] DFN-AAI Federation, see: http://www.dfn.de/content/dienstleistungen/dfnaai/ (German)

[5] Shibboleth Attribute Release Policy Editor SHARPE, see: J. R. Dalziel and Erik Vullings: *MAMS and middleware: The easily solved authentication, authorization, identity, single-sign-on, federation, trust, security, digital rights and automated access policy cluster of problems*. In: *Proceedings of EDUCAUSE 2005*.

[6] Wolfgang Hommel: *Using XACML for Privacy Control in SAML-based Identity Federations*. In: *Proceedings of the 9th Conference on Communications and Multimedia Security (CMS 2005)*, Salzburg, Austria, September 2005.

[7] Tim Moses (Editor): *OASIS eXtensible Access Control Markup Language 2.0, core specification*. OASIS XACML Technical Committee Standard, 2005.

[8] Matthias Ebert: *Konzeption und Implementierung einer policy-basierten Privacy Management Architektur für föderierte Identitätsmanagementsysteme am Beispiel Shibboleth*. Diploma Thesis, 2006 (German).

[9] Calvin Powers and Matthias Schunter: *Enterprise Privacy Authorization Language — EPAL 1.2*. Technical report RZ 3485 (#93951), IBM Research, Zurich, 2003.

[10] Seth Proctor: *Sun's XACML implementation*, see: http://sunxacml.sf.net/

[11] Matthias Ebert: Shibboleth XACML ARPs summary and download, see: https://spaces.internet2.edu/display/SHIB/ShibXACML

[12] ITIL.ORG: *ISO/IEC 20000 IT Service Management Standard*, see: http://www.itil.org/en/isoiec20000/index.php.

[13] IT Governance Institute and Office of Government Commerce, *Aligning CobiT, ITIL and ISO 17799 for Business Benefit – A Management Briefing from ITGI and OGC*, 2005.

[14] OTRS: Open Ticket Request System, see: http://otrs.org/.

[15] Maximilian Härtl: *Konzeption und Realisierung der technischen Unterstützung eines zentralen IT-Service-Desk mit OTRS an der TU München*, Diploma Thesis, to appear in 07-2007.

[16] Project elecTUM, see: http://portal.mytum.de/iuk/electum/newsboard

[17] Project mediaTUM, see: http://portal.mytum.de/iuk/integratum/bibliothek/index_html