

the information provided by CM about assets, software, incidents, known errors, changes, and releases, as well as data about staff, suppliers, locations, and much more. For example, the ITSM change management process determines the impact of requested changes on the infrastructure based on the CM data.

Built upon object-oriented principles, all the CM information is modeled as Configuration Items (CIs) on the one hand, and relationships – such as dependencies – between these CIs on the other hand. This data is then stored in a so-called Configuration Management Database (CMDB).

In its version 2, ITIL referred to the CMDB as an information nexus that had to be implemented by one single database [13]. Due to the vast amount of data that shall be stored in a CMDB, and due to the large number of potential sources of CI data, the cost and technical complexity for building one omniscient database were much too high for most enterprises to succeed. ITILv3, which is a major revision released in 2007, introduced the concept of a *CMDB federation* (CMDBf), which postulates that a holistic logical view should be built on top of an arbitrary distribution of CI data to existing databases and data repositories [15].

While ITIL defines the big picture and the ITSM processes, implementation details are left to ITSM tool vendors. HP and several other major vendors have founded a committee which works on the design of a CMDBf system and its interfaces [3]. In this paper, we present two extensions to the current CMDBf specification. The first is the extension of the CMDBf usage for inter-organizational services and second a policy based access control solution. Both concepts are described in the following.

1.1 Inter-organizational use of CMDBf

The selective sharing of CM information is an important aspect of outsourcing scenarios, for example when IT services and the IT service desk are operated by different companies for the same customer. Unfortunately, it has been ignored too long and was not adequately tool supported in the past. According to ITIL, the management of simple hierarchic customer-provider relationships as well as the handling of external IT service providers is performed by the discipline of Service Level Management. Within the scope of Service Level Agreements (SLAs) all service parameters are agreed to by customers and IT service providers. Concerning the operative part of customer-provider relationship management, ITIL recommends to establish interfaces between the existing CMDBs, both to the external and the internal IT service providers. This approach scales well when only a small number of service providers are working together to establish a service. However, this solution does not scale sufficiently in the case of complex multi-domain environments.

Throughout this paper, we will discuss the example shown in figure 1, which shows a very small subset of the real-world situation in the Munich scientific network (MWN): Most of the central IT services of the Technische Universität München (TUM), including their identity management system, e-mail services, and file and web servers are operated by the Leibniz Supercomputing Centre

(LRZ); this results in a very tight coupling of TUM's business processes with LRZ's obligations. However, TUM has its own service desk to support its staff and students, and so carrying out the incident management process efficiently requires access to up-to-date configuration management data across organizational borders. At first sight, this looks like there is a simple customer-provider relationship in place. This is why it has to be mentioned that every organizational unit of TUM – the picture shows only the physics department, but there are 11 further departments and also additional central institutes – acts as an individual IT service center having its own IT and being distributed on TUM's three major campuses in central Munich, Garching, and Weihenstephan. For such scenarios an inter-organizational CMDBf is urgently needed.

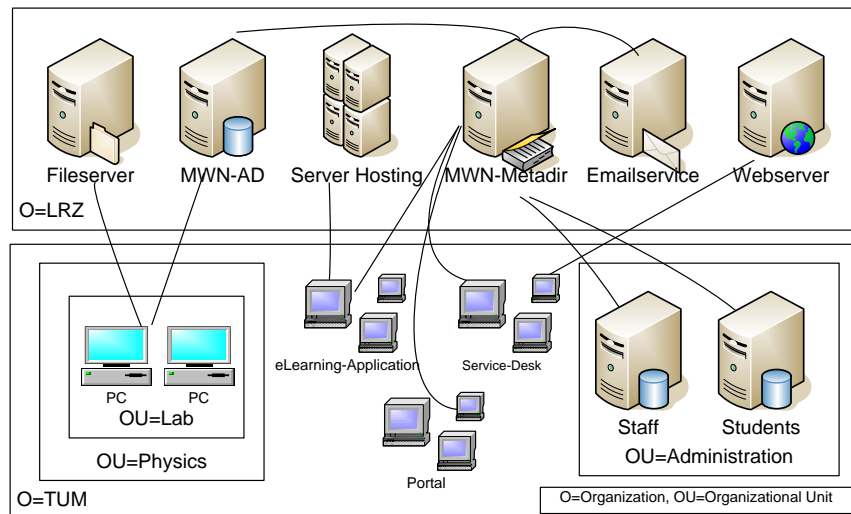


Fig. 1. Excerpt of the services operated in the Munich scientific network

While a CMDBf according to ITIL is not restricted to a single organization by definition, all ITIL processes have been specified with a single enterprise in mind; thus, we use the term ioCMDBf to emphasize the inter-organizational collaboration in providing the services.

Inter-organizational IT services share the following properties which are relevant for our work:

The whole is worth more than the parts: The inter-organizational IT service as used by the customers is provided through the collaboration of various independent IT service providers. In our scenario there are locally operated

directory services at the LRZ and a locally operated learning management application at TUM. But only via the collaboration of these a personalized access for students to e.g. register for exams is made possible. The IT support of the student's life cycle would be impossible without either of both services.

ITSM crosses organizational borders: Every organization is administered independently and changes also might be done independently in every organization. In our example there is no central control spanning LRZ and TUM. Every organization is managing its own resources like servers or applications. The inter-organizational services in turn are composed of such local services or resources. However, local ITSM is not sufficient any more. For example, if each organizational unit would only schedule its changes internally, the impact on collaboratively provided services could be huge. Thus, ITSM has to cross organizational borders, i. e. an inter-organizational ITSM needs to be established.

Figure 2 shows the relationship between the ioCMDBf and CMDBf in our example. The CMDBf federates the various management data repositories (MDRs) that are already in place in each organization to support the internal ITSM processes. For inter-organizational services it is evidently necessary, that also these CMDBfs are selectively logically merged into an ioCMDBf. The main difference to the local CMDBf is that each organization's scope and thus the areas of complete internal control are left. An important aspect for such a federation is access control, which we discuss next.

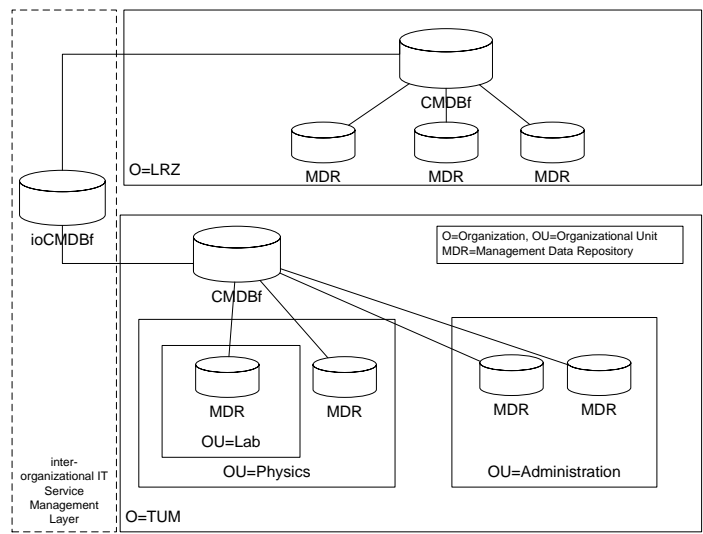


Fig. 2. ioCMDBf for inter-organizational Services

1.2 Access Control solution for ioCMDBf

We propose a policy-based access control solution for ioCMDBf data. While [3] deliberately does not cover security aspects of CMDBfs in depth, unauthorized access to an ioCMDBf obviously must be prevented, especially when more than two organizations are involved [4]. In order to reduce the administrative overhead for user management, we make use of federated identity management technologies.

We present our ioCMDBf concept and its access control management architecture in the next section. Its implementation in a real world project at TUM and LRZ has started just recently and is outlined in section 3. Competitive approaches and related work are discussed in section 4; a summary and an outlook to our next steps conclude this paper.

2 Solution: Inter-organizational use of ITIL federated Configuration Management databases

Sharing configuration management data across organizational boundaries enables more efficient IT service management, but also increases the complexity of the ITSM tool landscape. In this section, we first outline how an ioCMDBf can be used by the ITSM processes. We then discuss important aspects of the ioCMDBf information model in section 2.2; the access control mechanisms and workflows, which this paper focuses on, are presented in section 2.3.

2.1 Processes for ioCMDBf usage

To manage inter-organizational services efficiently the introduction of an inter-organizational ITSM (ioITSM) is necessary. The ioCMDBf supports the management processes of ioITSM like the CMDBf supports the management processes of ITSM. Management processes interacting with the ioCMDBf are on the line of the ITIL management processes, but they need to be enhanced for inter-organizational usage. Thus a couple of new roles and activities must be introduced. The advantage of orienting the ioITSM on ITIL is that a common understanding between the organizations arises at no cost. In this section we briefly describe processes that need to retrieve informations from the ioCMDBf.

Service Desk, Incident and Problem Management The service desk and the incident management process absolutely require an accurate overview of the composition of all services. The ioCMDBf assists in the error location task and provides the relevant information to delegate trouble tickets to the correct contact person even across organizational borders. Since it is not preferable that the users themselves have to look for the right service desk to contact, a new single point of contact, referred to as inter-organizational Service Desk, should be established [1]. The problem management tasks profit by the option to perform an impact analysis which takes inter-organizational relationships and dependencies into account. Having an ioCMDBf in place supports to implement partially automated processes for impact analysis [7].

Change Management Change management becomes very complex in multi-organizational cooperations. Sharing CM data allows to take dependencies across organizational borders into account and thus anticipate effects of planned changes more efficiently. An ioCMDBf supports the work of the necessary inter-organizational change advisory board (CAB) and should be complemented by a shared forward schedule of changes (FSC).

Release Management The release management process also greatly benefits from an ioCMDBf. In our scenario, the LRZ manages the majority of the software licenses acquired by TUM. Using an ioCMDBf, license information can be used by both sides without data redundancy or complex data synchronization processes. License usage information, which was previously only available to TUM, can then also be used by LRZ, which provides important knowledge for future negotiations with the software vendors.

Service Management The management processes of Service Design are also relevant consignees for the ioCMDBf [14]. In the case of Service Level Management in our multi domain szenario the definition of Service Level Agreements needs to be supported by information like service dependencies or supplier details.

Since the whole ITSM framework needs to be reconsidered for the use case of an ioCMDBf, a detailed role and process model will given in future work [12].

However, the intra-organizational compliance with the ITIL reference processes is a strong prerequisite for the use of an ioCMDBf. Only when all the involved organizations share a common understanding of the ITSM processes and use the same vocabulary, e. g. for peoples' roles, such as *configuration manager*, the ioCMDBf information model and access control will work.

2.2 ioCMDBf information model

An information model defines methods to model and describe managed objects (CIs) [8]. For brevity, we only outline the object relationships of our ioCMDBf data model as well as the query model here, but omit the guidelines for the detailed definition of CI attributes in concrete inter-organizational scenarios.

Figure 3 shows the data model with its basic elements: CIs may be simple or structured, and relationships can be arbitrarily parametrized; the profiles of contact persons are assigned to their respective organizations and ITIL roles.

Operations on CIs are triggered by ioCMDBf queries. Although we use a web services based interface, the query language can be compared, e. g. to SQL: Each query must indicate the type of the requested operation, i. e. creating, modifying, deleting, or searching one or more CIs. CIs are identified by object names; the ioCMDBf namespace includes an organizational prefix to prevent namespace clashes, which for example could occur when identically named objects of intra-organizational CMDBs are made available through an ioCMDBf. Conditions similar to SQL *where*-clauses can be used to restrict the number of affected CIs based on their attribute values.

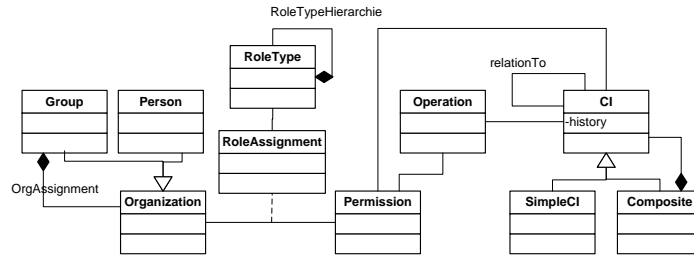


Fig. 3. ioCMDBf information model

2.3 Policy-based federated access control for ioCMDBfs

A CMDB federation stores a large number of CIs, relationships, and records; it also typically includes highly sensitive information, e. g. for financial management. Due to the large number of users which may access the CMDB through various ITSM tools and applications in the inter-organizational use case, access control is a crucial component of our ioCMDBf architecture. The underlying access control model is closely related to the auditing processes, which in turn are essential for governance, risk management, and compliance (GRC); clearly, the inter-organizational nature of our approach increases the solution complexity.

Our architecture is based upon three major design decisions:

1. The ioCMDBf does not have its own dedicated user management component. Instead, we rely on the Federated Identity Management protocol SAML, which allows us to retrieve the current user's profile from her home organization at run-time, including e.g. her name, email address, and roles, even if this profile has not been stored in the ioCMDBf a priori. This approach reduces both, the administrative overhead of managing external users and the risk of relying on outdated user information.
2. Our access management approach is policy-based and makes use of Attribute Based Access Control (ABAC). ABAC is a generalization of traditional role-based access control (RBAC), in which not only the user's roles are considered, but also the other attributes of the user object attributes, e. g. the user's department within her home organization. Access control rules are formulated as XACML policies, which we have successfully used in combination with identity federations previously [10].
3. By design, the access control we employ is very fine-grained. For CIs, it is applied on the attribute level, not just for a CI as a whole, contrarily to previous CMDB approaches such as [16]. Our intention is to support complex CMDB data models, in which CIs and records can have an arbitrary large number of attributes; obviously, sensitive attribute values must not be revealed to all users.

For user u and ioCMDBf object o , a policy p specifies whether action a may be performed under condition c . Conditions can make arbitrary use of

environmental data e , such as the current date and time, as well as user and ioCMDBf object attributes. Any given query q may affect a set of objects O and be affected by a set of policies P which use a set of conditions C . Thus, the result r of the policy evaluation f is $r_q = f(P, u, a, O, C, e)$ and takes the value of *Permit* or *Deny*.

As an example, *read* access to the *ip_address* attribute of a *server* CI can easily be restricted to users which have the role of *configuration manager* in the organization which the machine belongs to.

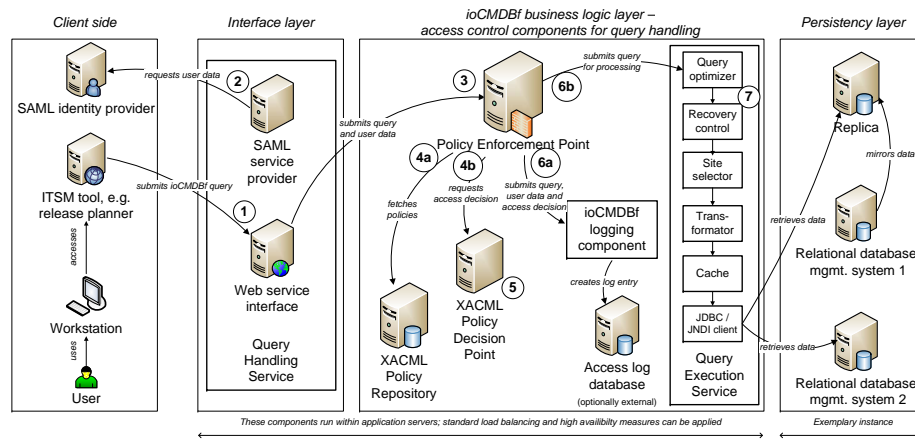


Fig. 4. ioCMDBf architecture with focus on access control

Figure 4 shows a simplified view of our ioCMDBf 3-tier architecture concept; it focuses on the access control components. The following workflow is executed for each request:

1. ITSM tools are ioCMDBf clients and use a web services (i.e., SOAP over HTTPS) based service access point, which is part of our query handler service (QHS) that is located in the ioCMDBf’s interface layer.
2. Queries to the ioCMDBf must either include a SAML attribute assertion, which provides the necessary user profile, or a reference to the user’s identity provider and the user’s id, which can in turn be used to request the user profile via SAML from the user’s home organization. For the latter use case, a SAML service provider component is integrated in the QHS.
3. The QHS extracts the user’s identity and attributes from the SAML assertion. This user data and the query are passed on to a policy enforcement point (PEP), which is located in the ioCMDBf’s business logic layer.
4. First, the PEP retrieves all the relevant XACML access control policies from the policy repository.³ The relevancy of a XACML policy can be determined

³ The ability to use distributed policy stores is an obvious requirement and will be part of our future work.

by matching its **Target** node set against the policy subject, action and object information derived from the query and the user profile.

Then, the PEP assembles a XACML request from the gathered user, query, and policy data in order to submit it to any XACML compliant policy decision point (PDP). In our scenario, a XACML PDP can be used that has already been deployed for the handling of privacy policies see ([9]).

5. The XACML PDP evaluates the request and returns the access control decision; for this functionality, no ioCMDBf specific extensions are required. XACML *Permit* and *Deny* decisions signal regular operations. XACML *NotApplicable* and *Indeterminate* decisions reflect errors due to missing or malformed policies; they are reported to the ioCMDBf administrators and are otherwise treated like a *Deny* decision.
6. The PEP receives the PDP's decision and passes it on to a logging component, which creates a new log database entry for each query. If the query was denied, the QHS returns an error to the web service caller. Otherwise, the query is sent to the query execution service (QES).
7. The QES can now safely assume that the user has the privileges required to perform the query. The details of the QES operations are outside the scope of this paper; however, it is important to note that for auditing purposes, modifications to ioCMDBf objects are also logged to their *history* attribute automatically.

The described architecture is complemented by policy administration points (PAPs); they are used to create, modify, and delete the XACML policies. Of course, the policies are subject to ITSM processes such as change and release management and are represented as CIs in the ioCMDBf.

3 Implementation in a real world project

The concepts presented in the previous section are currently being implemented as a part of TUM's service desk project [11]. This project has been started after it became evident that the recentralization of service operation necessitates likewise recentralized service support — even in inherently distributed and heterogeneous environments, such as German universities.

However, TUM's service desk is operated by TUM staff, although many of the services, including the identity management system, on which the trouble ticket system relies, are hosted by the LRZ. Half a year of TUM service desk operation has given us good insight into the current demands which Incident and Problem Management have towards the underlying CMDB. This allows a fine-grained design of CIs and their relationships. However, as modelling each of a whole university's CIs will be a very tedious task, we start with the CIs relevant for TUM's and LRZ's identity management, email, and e-learning systems; this service selection is based on the trouble ticket statistics and reflects the highest demands of service desk staff.

Our implementation will use TUM's trouble ticket system, for which the well-known open source software solution OTRS (cf. <http://otrs.org/>) is used and

which already offers some ITSM modules, as its primary management frontend. Written in Perl, OTRS is easily extensible to make web service calls to the ioCMDBf and make the Configuration Management data available to our first and second level support.

4 Related work and competitive approaches

Due to the inherent complexity of Configuration Management, suitable CM tools have shown to be hard to design; consequently, software suites intended to cover all ITSM processes often fail short of providing more than a very simple CMDB solution. Furthermore, most research focuses on process definitions, but not on tool support; especially the focus of our work, i. e. reference process based tool support for ITSM across organizational borders, is still a rather young discipline. In this section we summarize the current state of the art and its influence on our work.

CMDBs without their federation or inter-organizational aspects have hardly been investigated by research at all; often research has been done under the mere assumption that a suitable CMDB exists, without going into its details. Most ITSM tool vendors use a relational database management system for their CMDB, and the software's flexibility and usability depend on whether and how the design of customized CIs and relationships is supported. As of today, CMDB federation is the way to go. The ITIL standard itself proposes a CMDBf and even presents an architectural overview [15], without, however, discussing the details of how a CMDBf can be established in practice.

Consequently, the predominant understanding of this type of federation is that a CMDBf consists of one master database and several connected databases. The master database stores the so-called *core CIs*, while the connected databases store related information like incident records and service level agreements. Some products, such as [16], provide the capability to automatically link core and related CIs. Concerning research, [2] goes into some details about how CIs should be designed in the federated use case. However, several vendors still pursue the implementation of a single large database instead of a distributed CMDBf; for example, [5] resembles a virtual data warehouse approach.

Furthermore, most success reports about CMDBf implementations do not cover security aspects; for example, [17] and [6] assume closed environments. As a result, our work's contribution is twofold: First, our design has security (in terms of authentication, authorization, and auditing) in mind; second, our approach supports inter-organizational ITSM processes.

5 Summary of current status and next steps

In this paper, we have first motivated the need of using an inter-organizational configuration management database by means of a real-world scenario. We then outlined how ITIL-based IT service management processes will use our ioCMDBf and sketched the underlying information model regarding data storage and

querying. As this paper's focus, we presented the architecture and workflow of the policy-based ioCMDBf access control mechanism; taking security into account in the CMDB design phase is utterly important, even though the discussion of related work has shown that it often has been neglected.

The concept presented in this work is currently being implemented in a real world project, as discussed in section 3. Besides the incorporation of feedback gained from this deployment, we will refine the access control architecture, e.g. to support distributed policy repositories, as well as the auditing mechanisms based on additional requirements from the governance perspective. Concerning ioCMDBf usage, we will work on a detailed specification of the inter-organizational ITSM aspects based on the current ITIL reference processes. We also work on cookbook-style guidelines for the modelling of ioCMDBf CIs.

Acknowledgment The authors wish to 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/>.

The authors also 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/>).

References

1. Preeti Bhoj, Deborah Caswell, Sailesh Chutani, Gita Gopal, and Marta Kosarchyn. Management of new federated services. Technical Report HPL-96-131, HP Laboratories, December 1996.
2. Michael Brenner, Thomas Schaaf, and et. al. CMDB - Yet Another MIB? In *Proc. of 17th International Workshop on Distributed Systems: Operations and Management*. Springer Berlin / Heidelberg, 2006.
3. Forest Carlisle, Klaus Wurster, and et. al. CMDB Federation (CMDBf) - Committee Draft. Technical report, BMC Software, CA, Fujitsu, Hewlett-Packard, IBM, Microsoft, January 2008. <http://cmdbf.org/>.
4. Ronni J. Colville. Cmdb or configuration database: Know the difference, March 2006. Gartner RAS Core Research Note G00137125.
5. Troy Du Moulin. The Federated CMDB - Three Application of the Term, November 2006. Pink Elephant.
6. Denise Dubie. University taps ITIL to build open source CMDB. *Networkworld*, April 2007. Available online at <http://www.networkworld.com/>.
7. Andreas Hanemann, Martin Sailer, and David Schmitz. A Framework for Failure Impact Analysis and Recovery with Respect to Service Level Agreements. In IEEE, editor, *Proceedings of the IEEE International Conference on Services Computing (SCC 2005)*, Orlando, Florida, USA, Juli 2005.
8. Heinz-Gerd Hegering, Sebastian Abeck, and Bernhard Neumair. *Integrated Management of Networked Systems: Concepts, Architectures, and Their Operational Application*. Morgan Kaufmann, 1999.

9. Wolfgang Hommel. Using XACML for Privacy Control in SAML based Identity Federations. In *Proc. of 9th Conference on Communications and Multimedia Security*. Springer, September 2005.
10. Wolfgang Hommel. *Architektur- und Werkzeugkonzepte für föderiertes Identitäts-Management*. PhD thesis, Ludwig-Maximilians-Universität München, July 2007.
11. Wolfgang Hommel and Silvia Knittl. SERVUS@TUM: User-Centric IT Service Support and Privacy Management. In *Proceedings of 13th International Conference of European University Information Systems (EUNIS 2007)*, 2007.
12. Wolfgang Hommel and Silvia Knittl. An inter-organizational configuration management database as key enabler for future it service management processes. In *Submitted to eChallenges e-2008*, October 2008.
13. OGC (Office of Government Commerce), editor. *Service Support*. IT Infrastructure Library (ITIL). The Stationary Office, Norwich, UK, 2000.
14. OGC (Office of Government Commerce), editor. *Service Design*. IT Infrastructure Library (ITIL). The Stationary Office, London, UK, 2007.
15. OGC (Office of Government Commerce), editor. *Service Transition*. IT Infrastructure Library (ITIL). The Stationary Office, London, UK, 2007.
16. BMC Software. Federation and a CMDB. Available online at www.bmc.com, No. 59249, 2005. White Paper.
17. Andrea Stern. Reinvesting the IT dollar: From Fire Fighting to Quality Strategic Services. *EDUCAUSE Quarterly*, 24(3):8–14, 2001.

Biography



Wolfgang Hommel has a PhD in computer science from Ludwig Maximilians University, Munich, and heads the identity management team at the Leibniz Supercomputing Center. His current research focuses on IT security/privacy management in large distributed systems, including identity federations and Grids.



Silvia Knittl carries a diploma degree in informatics from Ludwig-Maximilians-Universität (LMU) Munich and holds the ITIL Service Manager certificate. She currently works at Technischen Universität München for the project IntegraTUM. Her research area is Configuration Management in federated environments.