

Federated and Service-Oriented Identity Management at a University

Frank Schell¹, Thorsten Höllrigl¹, Hannes Hartenstein¹

¹Steinbuch Centre for Computing, University of Karlsruhe (TH), Karlsruhe Institute of Technology, Karlsruhe, Germany

[frank.schell | thorsten.hoellrigl | hannes.hartenstein] @kit.edu

Keywords

Federation, User Provisioning, Federated Identity Management, Service-Oriented Architecture

1. EXECUTIVE SUMMARY

In this contribution we state the following thesis: *The concept of federation represents a promising way to ease the establishment and operation of organizational and technical issues of identity management at a university.* This concept fits well for most universities due to the fact that typically universities consist of ‘somewhat’ independent organizational units like library, computing center, administration and various faculties, with each having their own identity repository or even local identity management. We show two main advantages of this conceptual view of a university. On the one side the identity management can be build up successively in a step-by-step manner. On the other side the organizational units are seen as satellites with each needing just one or a small number of interfaces to the overall identity management system thus setting up a kind of hierarchy of identity management systems. This system can use different technologies, namely identity as a service and (de-)provisioning, to provide identity information to the organizational units and processes across the university. We exemplify how to integrate a satellite in the federation. Another contribution is the structuring of the establishment of a federation by categorizing artifacts and components in four models and by proposing a reasonable sequence of phases. This classification distinguishes between information, functional, communication and organizational aspects known from the integrated management of distributed systems. We conclude the paper with a discussion reflecting our experiences gained while setting up an identity management for a university.

2. Introduction

An integrated information management for students, scientists and employees at a university requires as fundamental building block an integrated identity management to support learning, teaching and research business processes. Thus, the question is not whether a university-wide identity management as the foundation of effective and efficient IT-services and of required IT-security is needed, but how to smoothly integrate it into the existing technological landscape and organizational structure. Based on our experiences in establishing such an integrated identity management we want to show the feasibility and the advantages when using the following thesis as a guiding idea:

The concept of federation presents a promising way to ease the establishment and operation of organizational and technical issues of Identity Management

Hereby we don't see a federation as a dogmatic and fundamentalist point of view or a legal form but as a helpful instrument. This concept fits well due to the fact that the organizational structure of common universities is resembled by 'realms' (or satellites). As 'realms' can be considered e.g. the computing center, university library, faculties and institutes. These organizational units evolved separately over the years, establishing their own IT-services, business processes and identity repositories. As a major benefit of the conceptual view "university as a federation" the various realms still keep the data sovereignty and so there is no need to replicate the identity information in a central system. Also, the access control for the respective data remains at the satellites and there is no need for centralized access control permissions. But two main advantages that we want to show in this contribution are that the concept allows the successive extension and therefore the integration of further organizational units in a step by step manner. So a federation can grow successively in scale and scope, from a few federation members to a large federation. The second advantage is that we achieve a kind of modular design with the infrastructural complexity of an organizational unit encapsulated in a black box manner through a minimum number of interfaces to the satellites.

3. Related Work

The term 'federation' is defined in different specifications and standards. WS-Federation [WSFed06] defines a federation as follows:

"A federation is a collection of realms that have established a producer-consumer relationship whereby one realm can provide authorized access to a resource it manages based on an identity, and possibly associated attributes, that are asserted in another realm. Federation requires trust such that a Relying Party can make a well-informed access control decision based on the credibility of identity and attribute data that is vouched for by another realm."

Another interesting definition is provided by [DD05]:

"A federation can be understood as a collection of trust realms that have established a degree of trust. The level of trust between them may vary."

Generally we understand a federation as a collection of independent organizational units that have established a trust relationship with each other.

Trust is therefore playing an important role in a federation. The Liberty Alliance [Linn03], which is a consortium of about 150 organizations and companies with the goal to create standardized federation protocols without disregarding privacy, is distinguishing between pair wise, brokered and community trust depending on the business model. Based on these trust models it is possible to define three types of federations as it is likewise done in [Wind05]: An *Ad-hoc federation*, a *Hub-and-Spoke federation* and an *identity federation network*. An Ad-hoc federation is built on bilateral agreements of the respective federation members. That means a new federation partner has to establish a business agreement to every federation partner he wants to do business with. In the worst case this means that every federation partner has to establish and manage a trust relationship with every other federation partner. In a Hub-and-Spoke federation this drawback gets eliminated by having a centralized component. A new federation partner only has to establish a trust relationship

to the central component which deals as a trust broker. The complexity for the needed trust relationships is therefore massively reduced. A drawback for this kind of federation is the central component, mostly one of the larger federation partners, which can overrule smaller federation partners. To avoid this, an identity federation network is setting up a delegation of representatives of the federation members which are deciding about federation relevant issues. The affiliation to an identity federation network is achieved by becoming a member of this 'community' and so the trust relationships between the federation members are built through the trust to the community.

The choice of the right kind of federation for an organization depends on different factors, like the size of the university, the degree of decentralization, the desire for autarky of the single institutions or the potential rate of expansion, just to mention a few. In a common university we suggest to establish an identity federation network due to the fact that the organizational units of a university are already members of a community represented through the university.

4. Architectural and Interface Aspects

In a federation almost all satellites, processes across organizational units, portals for user interaction, and IT-Services in general need identity information. We see two different ways how to provide this information, namely 'Identity as a Service' (IDaaS), and user provisioning. We will discuss the pros and cons of each approach by looking at the interfaces, and the ease of their federation-wide integration.

First, this data can be provided by services spread all over the federation. This fits perfectly in our view of the university as a federation, because the identity information is provided by the satellites which know best how to manage and who is allowed to access this data. Therefore the concept of federation supports the advancement of new services by leaving the control of the identity information and consequently the provided services by the satellites. Thus the acceptance for Identity as a Service increases through the concept of federation. We identified two different types of services in a federation. On the one hand there are several attribute services (attribute authorities) that are encapsulating legacy systems of satellites, e.g. HR systems, through standardized interfaces, e.g. Web Services. On the other hand there are several infrastructure services, namely authentication service, authorization service, identifier mapping service or role mapping service. The identifier mapping service is the fundamental part to link the different accounts of the diverse organizational units. A member of a university normally has multiple accounts spread all over the variable institutions. These accounts do not have any connection to each other. So it is typically not possible to check e-mails at the computing center and then check the library account without authenticating at different systems. The identifier mapping service delivers the needed connection information and therefore allows university-wide processes. Authentication services have the task to authenticate users possibly in different ways, like passwords or user certificates, and vouch for this authentication. Authorization services can make authorization decisions based on different aspects, e.g. the type of authentication, the properties or roles of a user. A role mapping service acts as intermediary between different realms for exchanging roles defined in one realm to a similar role of another realm. Even if this is the most preferred way to provide identity information not all legacy systems can invoke services to obtain this data. Another problem here is the potential increasing number of service calls that have to be made to gather all information from various attribute authorities, necessary e.g. for authorization decisions. However, one of the most reasonable advantages of IDaaS is the reuse of the attribute services and infrastructure services in different contexts. Various federation-wide business processes and applications can invoke or orchestrate these services to reduce overhead and to increase efficiency of the federation.

Another solution is a user provisioning system for providing identity information. This is necessary, because there are legacy systems which must have a local identity repository and are not able to get this information from a remote system. Thus it appears that there is the need for a provisioning system which can deliver identity information from authoritative sources to target resources throughout the federation without the need of storing this information in a central directory. The provisioning processes can be actively triggered through defined interfaces of the provisioning system to push actual data at any time into the target resources. As long as there is data redundancy at the organizational units there is a need for such a provisioning system to achieve federation-wide data consistency, even in a Service-Oriented Architecture where different services provide the same

data such a mechanism is needed. Though such a provisioning system can achieve data consistency across the university this system is increasing the complexity of the identity management system. The concept of federation is decreasing this complexity through the decent of the connected systems by encapsulating a satellite by a minimum of interfaces.

Our Approach

The system we have set up is a combination of the former described concepts. That means we have established different services as attribute authorities as well as a provisioning system to push data to the satellites if applications or business processes can't get the needed identity information over the provided interfaces. As IDaaS e.g. a Web Service, called PersonService, encapsulates the organizational unit 'central administration' which uses the 'Hochschul-Informationssystem' - the most common software used by university administrations in Germany [Hi07]. This system has different databases which store most of the relevant identity-related data in a common German university. Therefore this service is the authoritative system for most identity information needed at other organizational units. This is the reason why we need to connect this service with our identity management system to provision identity information to other satellites and achieve university-wide data consistency if this data changes. To use this service we developed an adapter based on Web Service Interoperability Technology (WSIT) that uses common Web Service technology and WS*-specific protocols to securely access different Web Services [We07]. To enable actual security features we used the .NET implementation of the WS*-specifications, called Windows Communication Foundation (WCF) for realizing the Web Service [WCF07]. The service can be used to retrieve the necessary data in a standardized way and to transmit this data to the user provisioning system which is realized with the Sun Identity Manager (SIDM) as the core of the system [Su07]. We chose the Service Provisioning Markup Language (SPML) [Ro03], which provides an XML-based framework for managing the allocation of system resources within and between organizations, for transporting the provisioning information into the satellites. The SIDM can be easily modified to enable different provisioning processes by specifying them in XPRESS, an XML-based expression and scripting language. The SIDM can connect to further target resources through specific adapters that encapsulate the concrete technology of the resource. Figure 1 shows an overview of the technical components of our approach.




Figure 1 - Overview of the technical Components of our Approach

As a result of seeing the university as a federation of their organizational units each satellite can now be easily integrated, due to the fact that integrating a whole satellite means understanding just one single interface, namely a self-developed Web Service. There is no need to integrate the whole bunch of local repositories, like databases, directory services, and HR Systems, of this satellite directly in the identity management system. As a consequence the overall identity management system does not need to have any knowledge about the internal infrastructure of the satellite.

5. Procedure Model

To ease the launch of a federation we suggest a structuring of the establishment of a federation by categorizing artifacts and components in four models and by proposing a reasonable sequence of phases. Due to the fact that the establishment of a federation is as mentioned above an iterative process these phases must be seen as initial efforts. The advantage of our concept here is that it is not necessary to decide and define all relevant issues initially, but the artifacts can be adapted and extended at any time. Our classification distinguishes between information, functional, communication and organizational aspects known from the integrated management of distributed systems [ITUT97]. In figure 2 we show a division of the most important federation artifacts in these four models. The figure furthermore depicts the dependencies between the artifacts and components.

Organizational Model - Federation Fundamentals

The organization model supports different organizational aspects like fundamental agreements that have to be specified at the beginning of the establishment of the federation. These fundamental aspects are needed for the federation to be functionable and capable of acting. Without these decisions it is not possible to define further technical and organizational details. So first of all the federation should decide on organizational issues, like the type of a federation, decision procedure, and naming and definitions.

- Type of federation: A federation has to agree on a type of federation: ad-hoc, hub-and-spoke, identity federation network. This agreement has a fundamental impact on all other organizational and technical issues.
- Decision procedure: Depending on the type of the federation there are many decisions that have to be made about all kinds of aspects, e.g. integration of new federation members, the information model, used security mechanisms and so on. The federation has to establish an appropriate decision procedure to make these decisions.
- Naming and definitions: The federation needs a consistent and uniform language to avoid misunderstandings and to achieve an efficient communication between the federation partners. An adequate method to achieve this is to arrange and maintain a glossary in form of a wiki or something similar.

Information Model

- The information model specifies the provided attributes of the federation and their associated meta data, like policies, attributes or service catalogue. These specifications deliver the syntax and semantics for federation-wide artifacts. This is an essential step for a common understanding of all relevant, distributed information. For example SAML specifies therefore three types of assertions, attribute, authentication, and authorization decision, to distribute this kind of data [Ca05].




Figure 2 - Categorization and Dependencies of Federation Artifacts and Components

Communication Model

The next step is to define the communication model, which describes the communication and transport protocols needed for the exchange of identity information. In a federation there can be used a variety of federated identity management protocols, like SAML, WS-Federation, and LA ID-WSF.

Functional Model

A federation has two kinds of services, each having a service description about the provided functionality. There are infrastructural services, like authentication, authorization, identifier mapping, and role mapping services and attribute authorities, like the previously described PersonService. Besides the general functionality the functional model specifies additional requirements for services, like fault handling, logging, monitoring, auditing, accounting, and performance management. To simplify the integration of services respectively their needed protection concerning the access and the communication it is useful to categorize the offered services such as it is done in the Liberty Identity Assurance Framework [Cutl07]. In this specification the Liberty Alliance suggests to divide services in four different levels depending on their strength of assured security.

Organizational Model - Single Member

Consequently after all the previous models are specified the federation has an essential common understanding about the provided services and attributes. Based on this information every satellite can describe its provided attributes and services in manner all the other federation members understand.

6. Conclusion

In this chapter we discuss the experiences made while establishing an identity management at our university. We consider the step-by-step establishment and the satellite-concept of a federation as prominent advantages. Through this it was possible to quickly achieve results, a fact that is important for the acceptance of an identity management project. Hence it was possible to reduce the organizational overhead by having less organizational meetings with less people and therefore coming to quicker results without losing the possibility to extend and adapt the established system. To view the satellites in a black box manner leads on the one side to less organizational tasks for the federation and on the other side encapsulates the satellite and therefore changes of the internal satellite systems. That means that e.g. the change of the HR System of an organizational unit has only impact on the interface and not on the overall identity management system and - what is even more important - no impact for the other federation partners. The responsible people for these interfaces can act as intermediaries between the federation and the internal organizational unit. This reduced the organizational overhead especially in the beginning of the project. Therefore, the identity management could go online without the involvement of all potential organizational units in advance. After the initial steps are made we can now integrate further satellites with little effort. The integration of new satellites leads, however, to challenges like the identifier and role mapping information, which have to be gathered and linked.

7. REFERENCES

- [Ca05] Cantor, S. et al.: Assertions and Protocol for the OASIS Security Assertion Markup Language (SAML) V2.0, OASIS Standard saml-core-2.0-os, march 2005
- [Cutl07] R. Cutler. Liberty Identity Assurance Framework 1.0, http://www.projectliberty.org/liberty/content/download/3736/24651/_le/liberty-identity-assurance-framework-v1.0.pdf, 2007. Liberty Alliance Project.
- [DD05] Djordjevic, I, Dimitrakos, T.: A note on the anatomy of federation. In: BT Technology Journal, Vol 23 No 4, october 2005
- [Hi07] HIS Hochschul-Informationen-System GmbH: <http://www.his.de/>, 2007
- [ITUT97] ITU-T - X.701: Information technology - Open Systems Interconnection - Systems Management Overview, <http://www.itu.int/rec/T-REC-X.701>, 1997.
- [Linn03] J. Linn. Liberty Trust Models Guidelines 1.0, http://www.projectliberty.org/liberty/content/download/1232/8000/_le/liberty-trust-models-guidelines-v1.0.pdf, 2003. Liberty Alliance Project.
- [Ro03] Rolls, D.: Service Provisioning Markup Language (SPML) Version 1.0, OASIS Standard, october 2003.
- [Su07] Sun Microsystems: Sun Java System Identity Manager. http://www.sun.com/software/products/identity_mgr/index.jsp, 2007
- [WCF07] Windows Communication Foundation of the Microsoft .net Framework 3.0. <http://wcf.netfx3.com>, 2007
- [We07] Web Service Interoperability Technology. <http://java.sun.com/webservices/interop/>, 2007
- [Wind05] P. J. Windley. Digital Identity. O'Reilly. 1. edition, 2005.
- [WSFed06] H. Lockhart, S. Andersen, J. Bohren, Y. Sverdlov, M. Hondo, et al. Web Services Federation Language 1.1 (WS-Federation), <http://specs.xmlsoap.org/ws/2006/12/federation/ws-federation.pdf>, december 2006.