Géant-TrustBroker: Simplifying Identity & Access Management for International Research Projects and Higher Education Communities

Daniela Pöhn[1], Stefan Metzger[1], Wolfgang Hommel[1]

[1]Leibniz Supercomputing Centre, Bavarian Academy of Sciences and Humanities,
Boltzmannstraße 1, D-85748 Garching n. Munich, Germany
{poehn,metzger,hommel}@lrz.de

Keywords
Federated Identity Management, Inter-Federation, Shibboleth, SAML, Trust management.

1. Abstract

Most national research and education networks (NRENs) have set up authentication and authorization infrastructures (AAIs), also known as federations, to ensure that ICT services can be used across higher education institutions' (HEIs') borders. For example, the German federation DFN-AAI allows students from various universities to enroll in eLearning courses provided by other German universities (Hommel, 2009). Most European federations are technically based on the SAML standard and implemented using open source software like Shibboleth or simpleSAMLphp. However, users can only access third party services whose service providers (SPs) are members of the same federation as their home organization, which is also referred to as their identity provider (IDP). Therefor, given national federations, international groups of users, e.g., researchers in a multi-national EC-funded project, cannot access each others' ICT services, such as a project-wide Wiki collaboration web server, without additional efforts simply because crossing federation borders is not possible technically.

In the past, many HEI members with a demand for international identity & access management (I&AM) have often worked around this problem in one of two less elegant ways: They either created local user accounts for their external project partners at each service, which does not scale well, or they created community-specific new federations, which were not defined by geographical but by any other arbitrary criteria, such as membership in a scientific community or project. However, neither of these solutions are user- and administration-friendly, but instead increase the overall management complexity and are considered burdensome overhead. To overcome the limits imposed by national federations, the pan-European research and education network Géant meanwhile initiated eduGAIN (see (Géant, 2014)), which is an umbrella inter-federation (i.e., a federation-of-federations) that enables Inter-AAI user authentication and authorization (AuthNZ). More than 20 federations world-wide already have joined eduGAIN, making it one of the most important eScience-enabling software infrastructures as of today.

eduGAIN, however, comes at the price of increased contractual complexity, and, on the technical side, has only standardized the common denominator of its federation members regarding which information about users IDPs make available to SPs. In practice, this means that there is no guarantee that users from an IDP in federation A can successfully use a service provided by an SP in federation B, even if both of them are in eduGAIN, in the same way as if the IDP and the SP were in the same (national) federation. Thus, while eduGAIN is certainly a success and enables the use of many services across federations' borders, its adoption turned out to process slower than initially hoped for and the created inter-federation by itself is not completely sufficient for more complex services that need more detailed user information from IDPs.

Géant has therefore initiated a project complementary to eduGAIN: Géant-TrustBroker (GNTB) will enable the on-demand creation of virtual federations and put the end users in control of connecting arbitrary SPs to their own IDP even when they are not in the same federation or eduGAIN. GNTB

optionally supports the fully automated setup of technical SP-IDP relationships so that users can immediately start using new services provided by federation-external SPs instead of having to wait until the SP and IDP administrators have set up the AAI software configuration manually. Manual intervention is only necessary when organizational trust-building measures, such as signing a formal contract between SP and IDP, are necessary, e.g., for commercial services that require high liability.

In this article, we present the concepts of GNTB from the perspective of a HEI that operates an IDP for its users, assuming that the IDP already is a member of at least one federation, typically the national NREN's AAI. We first discuss the motivation for GNTB from both the end users' and the HEIs' perspectives and show how GNTB can be used stand-alone or in conjunction with eduGAIN. We then give an overview over the functionality and technical workflows that GNTB implements, again with a focus on the IDP side. GNTB is currently being developed in Géant's GN3plus project and will be available for pilot use in 2015; we therefore conclude this article with a summary of what has been achieved so far and an outlook to our ongoing work.

2. MOTIVATION FOR GÉANT-TRUSTBROKER

Setting up an IDP software like Shibboleth for the first time is a challenge in itself for many system administrators, e.g., at a HEI's data center. It usually involves connecting the IDP software to a local I&AM system, such as an LDAP server, that can be used to authenticate one's users and to fetch user attributes, such as the user's email address or language preferences, which are required by services to be properly used. While this IDP-internal setup of the local I&AM interfaces needs to be performed only once, configuration gets more complicated when the IDP's users start to actively use external services: Given the technical nature of IDPs and SPs, the software on each side must have some information about the other side before the IDP can deliver user information to the SP. For example, the IDP needs to know the SP's communication endpoints (URLs) and the SP needs to know the IDP's server certificates that are used to digitally sign and encrypt the user information; this set of information required by the other side is commonly referred to as (IDP and SP) metadata.

As a consequence, IDP administrators must ensure that they have the metadata of each SP available that is used by at least one of their users. To avoid the overhead of adding each SP's metadata to the IDP's configuration manually, IDPs usually join their NREN's national federation. Federations, which are sets of logically grouped IDPs and SPs, provide several advantages over manually managing inter-organizational service access:

- There are contracts to ensure a minimum of desired behavior. For example, on the one hand IDPs in the federation must ensure that they provide high-quality user information, e.g., to avoid service misuse due to fake accounts; on the other hand, SPs must commit themselves to honor national or federation-specific privacy and data protection regulations so that IDPs can safely assume that their users' personal data will not be abused by the SPs.
- Federations aggregate the metadata of all member IDPs and SPs and make them available as the so-called federation metadata. Thus, by joining one federation, IDPs get the metadata of many SPs and vice versa.

On the IDP side, problems arise whenever there is a user locally at the HEI who wants to use an external SP that is not part of the same federation as the IDP: Because SP and IDP do not know each other, the user cannot login to the service directly. Up to now, this problem can only be solved as follows:

1. The IDP administrators manually add this SP's metadata to their configuration; this must also done by the SP administrators regarding the IDP metadata.
2. The IDP additionally joins any of the federations the SP is already a member of (or, less likely in practice, vice versa). This means that the organizational and technical efforts known from joining the national federation have to be repeated.
3. Both the IDP and the SP join an umbrella federation such as eduGAIN.

It is quite obvious that the first of these three approaches does not scale very well; the task of setting up IDP-SP connections manually is repetitive and tedious, and since metadata has a limited validity period, e.g., due to the included server certificates, must be repeated whenever the metadata changes. Similarly, joining an additional federation and maintaining membership status is like cracking a nut with a sledgehammer when there is only a small number of local users who actually want to use a single external SP - the technical result does not justify the effort. Having both sides join eduGAIN is the best solution so far but has the issues mentioned in the previous section; furthermore, as a federation-of-federations, eduGAIN is no solution for new stand-alone SPs, i.e., if the SP is not already member of another federation, the eduGAIN approach will not work either.

One important aspect of either of the three approaches is that the user has to wait until the IDP and SP administrators decided on a solution and finished implementing it. Currently, the user who wants to use a new service at some SP will see that the SP supports federated login, but her IDP is not in the list of organizations who can access the service this way. The user will then have to notify her IDP's administrators about the issue and wait until it has been solved. The currently available options therefore cause the following primary problems:

1. IDP administrators manually have to adapt the IDP's configuration to the new SP and ensure that the SP's metadata is sustainably kept up-to-date, which causes additional workload for each additional SP.
2. Users cannot immediately start to use a new service but have to wait until the SP and IDP administrators' manual work is done, which often does not meet the users' expectations and causes frustration, often leading to the situation that the user is no longer interested in this particular service or finds other ways to use it, e.g., by creating a local user account there.

In practice, service access across federation boundaries additionally often faces the problem of incompatible user data schemas. For example, one federation may specify that the user's real name should be provided by means of two attributes, e.g., givenName and surname, while the other federation uses a single attribute for the same purpose, e.g., fullName. Accessing services properly therefore will not work if the SP expects the attribute fullName while the IDP can only provide givenName and surname. To solve this problem, IDP software packages allow for the definition of user data conversion rules, i.e., for this particular SP, the IDP will assemble the fullName attribute out of its local givenName and surname attributes on-the-fly. This has the advantage that the IDP's local I&AM system does not need to be extended by an additional fullName attribute, but again, the user cannot properly access the service before the IDP administrator has implemented such a conversion rule.

The sum of these issues motivates a new approach for inter-federation service access with a focus on automation: By eliminating or at least greatly reducing the amount of manual configuration and implementation work that needs to be done, IDP and SP administrators can be relieved from awful routine tasks and users can access new external services much faster, ideally instantly. The GNTB project therefore develops the following solution:

- A GNTB service enables the user-triggered, on-demand exchange of IDP and SP metadata whenever the first user from the IDP tries to access the SP.
- A GNTB plug-in for the IDP software automatically reconfigures the IDP to set up user attribute release filters for the SP.
- User attribute conversion rules implemented by one IDP can be shared via a GNTB rule repository and re-used by other IDPs. Since IDPs in the same (national) federation use the same user data schema, it is usually sufficient when one IDP in the federation manually implements the required conversion rules, sparing other IDP administrators in the same federation the implementation efforts.

This enables the full automation of the IDP-SP communication setup procedure and grants users immediate full access to the service, at least in the majority of use cases when there are no user data conversion rules required or data conversion rules have already been shared by other IDPs via GNTB. Also, if full automation is not desired, because, for example, IDP administrators want to stay in full control over changes made to their IDP's configuration, manual approval steps are possible while still granting the comfort of automatically creating the necessary technical configuration changes. GNTB, however, does not address organizational aspects such as written contracts between IDPs and SPs; those must either be handled as before, i.e., individually for each SP, or can be addressed by means of joining an inter-federation like eduGAIN.

3. GÉANT-TRUSTBROKER FUNCTIONALITY

GNTB can be regarded as a software service for the provisioning of dynamic virtual federations: IDPs and SPs, which want to use GNTB functionality, first have to register at GNTB and upload their metadata, similarly to how they have to manage their metadata in other federations they join, except that registration is possible for any organization and does not require a written contract. However, unlike eduGAIN, GNTB does not bundle all the IDP and SP metadata entries to one huge inter-federation metadata file, but instead it makes single metadata entries available to other organizations on demand. In practice, this is a key difference because complete inter-federation metadata sets are huge XML files whose processing can noticeably slow down IDPs and SPs depending on the hardware and software that is used, and slower responses from IDPs and SPs immediately result in degraded usability from the end users' perspective. Instead, the bilateral metadata exchange triggered by GNTB results in a virtual federation starting with only two members, the SP and the IDP, and GNTB can be used to dynamically set up an arbitrary number of such virtual federations:

- First the user visits a service site and chooses his account, e.g. his institutional ID, similar to OpenID's accountchooser.com feature and the discovery service known from most NREN federations. This functionality allows those users, who have multiple accounts at different IDPs, e.g. because they belong to multiple organizations, to choose which of their digital identities they want to use for accessing the SP. As shown in Figure 1 GNTB could also be part of this IDP list
- In the next step the SP checks his trust relationship towards the user's IDP. If they already technically trust each other, then a normal (Inter-)FIM workflow applies without the involvement of GNTB. Otherwise, the GNTB trust establishment workflow will be triggered. Thus GNTB triggers the IDP and redirects the user for authentication to the login page of his identity provider in order to prevent Denial-of-Service attacks.
- When a session is successfully established, the IDP starts fetching metadata and attribute information about the SP, who has to register at the GNTB beforehand, from the Géant-TrustBroker.
- An important part of the GNTB core workflow is that it triggers the automated re-configuration of the IDP software if the IDP administrators make use of the GNTB IDP plug-in. Besides integrating the SP's metadata in the local configuration, this sub-workflow also automates the download and activation of user data conversion rules if they are required and suitable rules have already been shared by other IDP administrators, and re-configures the IDP to send only user attributes to the SP that are actually required by the service, which is also known as attribute filtering – otherwise, the IDP would send the complete information it has about the user to the SP, which would not comply with privacy protection principles such as data minimization and therefore violate, for example, the EU's data protection directive. We explain how IDP administrators can make use of this GNTB feature in section 5.
- After attribute filters are created, the assertion containing the user's attributes is digitally signed by the IDP and after the user's approval, e.g. by uApprove, sent to the SP in order to let the user access to the service.
- For verification purposes of this digital signature the SP now has to fetch the IDP's metadata information, i.e. certificates, from GNTB.
- Registering and keeping IDP metadata up-to-date is part of the so-called GNTB IDP management workflow, which is described in the next section.

The metadata exchange between an SP and an IDP with no previous relationship is triggered by the user who wants to access a service and referred to as the GNTB core workflow; it basically works by having the GNTB service act as a virtual IDP towards the SP and as a virtual SP towards the user's IDP. However, GNTB is only involved during the first-time setup of the SP-IDP communication; any subsequent logins to the service by the same and other users of the IDP work directly, so that GNTB cannot become a performance bottleneck, single point of failure, or raise any new privacy issues.

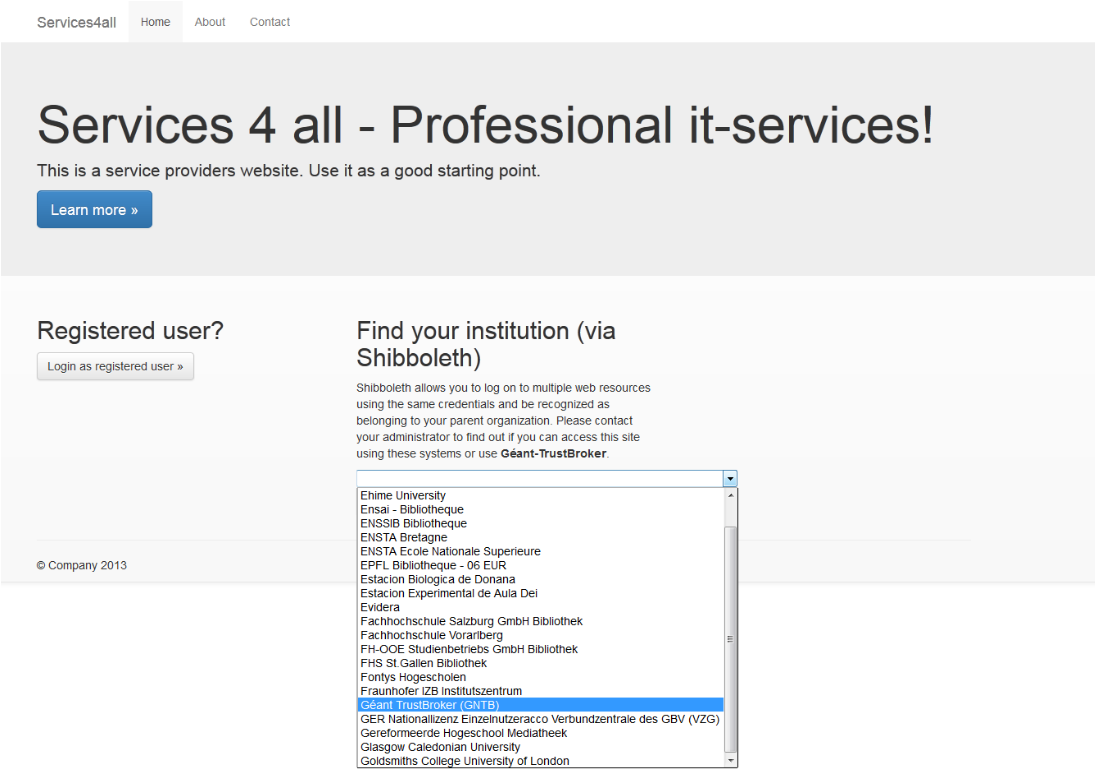

Services 4 all - Professional it-services!

This is a service providers website. Use it as a good starting point.

Registered user?

Find your institution (via Shibboleth)

Shibboleth allows you to log on to multiple web resources using the same credentials and be recognized as belonging to your parent organization. Please contact your administrator to find out if you can access this site using these systems or use **Géant-TrustBroker**.

Ehime University
Ensai - Bibliotheque
ENSSIB Bibliotheque
ENSTA Bretagne
ENSTA Ecole Nationale Superieure
EPFL Bibliotheque - 06 EUR
Estacion Biologica de Donana
Estacion Experimental de Aula Dei
Evidera
Fachhochschule Salzburg GmbH Bibliothek
Fachhochschule Vorarlberg
FH-OOE Studienbetriebs GmbH Bibliothek
FHS St.Gallen Bibliothek
Fontys Hogescholen
Fraunhofer IZB Institutszentrum
Géant TrustBroker (GNTB)
GER Nationallizenz Einzelnutzeracco Verbundzentrale des GBV (VZG)
Gereformeerde Hogeschool Mediatheek
Glasgow Caledonian University
Goldsmiths College University of London

Figure 1: Account choosing functionality

4. GÉANT-TRUSTBROKER IDP MANAGEMENT WORKFLOW

IDPs have to digitally sign the user information, called assertions, they send to the SPs in order to protect the communication. Therefore the SPs need access to the IDPs' certificates, which are part of their metadata, to verify the signature. Géant-TrustBroker helps both parties, while they manage their metadata information at the central Géant-TrustBroker.

The basic workflow of the IDP management is shown in Figure 2. After the initial metadata creation, which is also part of the intra-federation communication, the IDP metadata has to be registered during the first time contact between IDP and GNTB. As a precondition an authentication of IDP administrators is required, e.g. by providing a valid username/password, in order to prevent malicious intent. The metadata entries themselves are stored as signed XML files in a repository, which also provides version control functions, e.g. comparing two different versions and restore a previous one. Additionally, in an associated relational database GNTB stores required information like file system path, version number, and timestamp information for validation.

The IDP administrator has also possibilities to update the registered metadata if needed by uploading a complete new metadata file. If an IDP decides to no longer use GNTB, its metadata and also all provider-related information stored in the database has to be deleted. The same applies for SP administrators and their metadata.

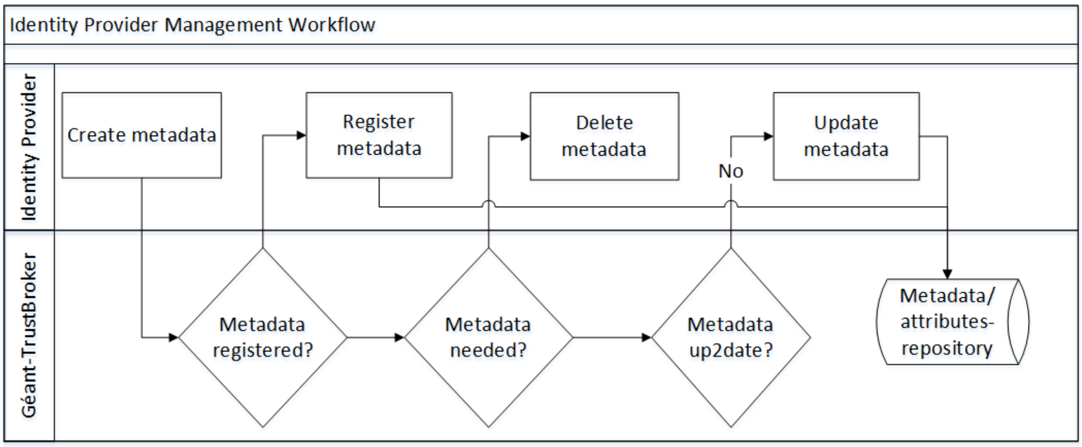

Figure 2: IDP Management Workflow

To support IDS administrators' daily tasks, the GNTB API provides a lot of functionalities to handle metadata management with command line tools or, more in a more manual way, via a web frontend. For example, IDP administrators can register their entity by file or URL. Additionally they should provide the type of entity, i.e. SP or IDP:

- gntb Ent_RegisterByFile(File, Type) -> ok, error
- gntb Ent_RegisterByUrl(Url, Type) -> ok, error

During trust establishment between an IDP and an SP, the IDP receives the SP's metadata via GNTB. Afterwards IDPs and SPs exchange requested user information and attributes directly without involving GNTB. If an SP updates its metadata including attribute information at GNTB, all IDPs, which already have a technical trust relationship with this specific SP, are informed about the update or fetch this new data automatically, e.g., pushed by GNTB or periodically pulled by the IDP. This workflow is shown in Figure 3.

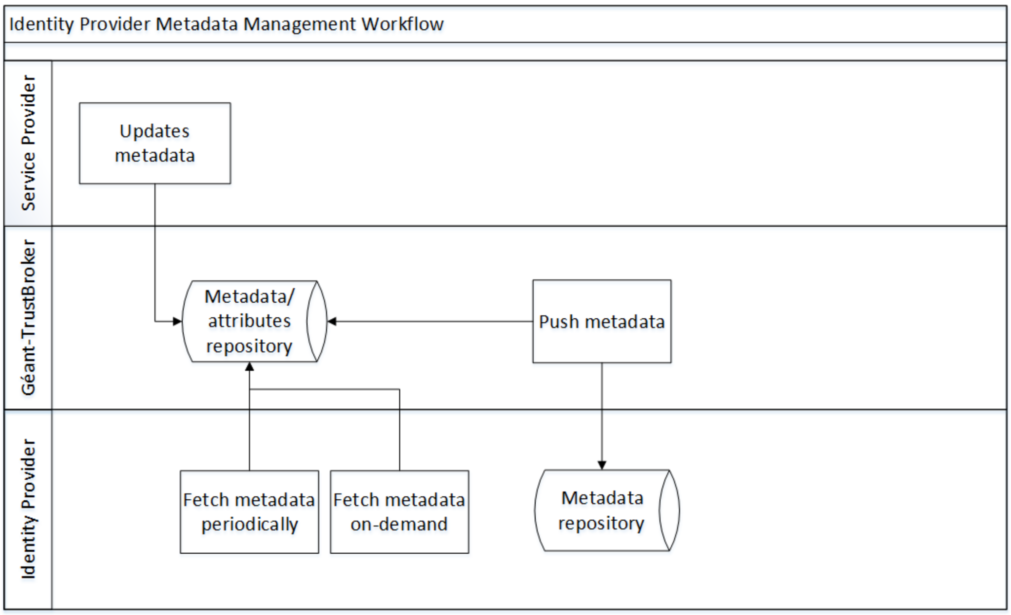

Figure 3: IDP Metadata Management Workflow

IDP administrators can set two different notification types:

- gntb Conv_SetNotify(Rule_ID/Name, Boolean) -> ok, error: If set true, administrators are going to receive notifications about changed conversion rules, which they have reused in the past.
- gntb Ent_Notify(Entity_ID, Boolean) -> ok, error: If set true, administrators are going to be notified, if metadata of technical trusted SPs is updated.

The notification is sent usually by e-mail, but updates can also be fetched automatically or pulled periodically, which can be configured by parameters on the IDP side as a normal Metadataprovider.

5. GÉANT-TRUSTBROKER CONVERSION RULE MANAGEMENT

The biggest waiting time for users occurs when SPs want more attributes than those that are delivered in the default IDP configuration. Those attributes may not even be part of the IDP's schema, i.e., IDPs cannot provide them out of the box. IDP administrators create these new attributes or map existing ones to SP-specific attributes by writing so called conversion rules. In order to search, share, and re-use them with administrators of other IDP that utilize the same schema, e.g. because they are members of the same federation, these conversion rules could be uploaded to the GNTB conversion rules repository, as shown in Figure 4. When conversion rules already exist, the IDP administrators get a notification and have the possibility to download them or even have them integrated into their IDP configuration automatically. IDP administrators have read access to all conversion rules, but only write access to their own rules, e.g. to delete or update them. If an SP updates one of its services, which needs from now on more attributes, then also conversion rules could be affected and existing rulesets are flagged as outdated. Based on this flag those administrators, which reuse these rules, get a notification, if the activate the SetRuleNotify option or have the possibility to update and activate them automatically in an auto-update function once a new version becomes available.

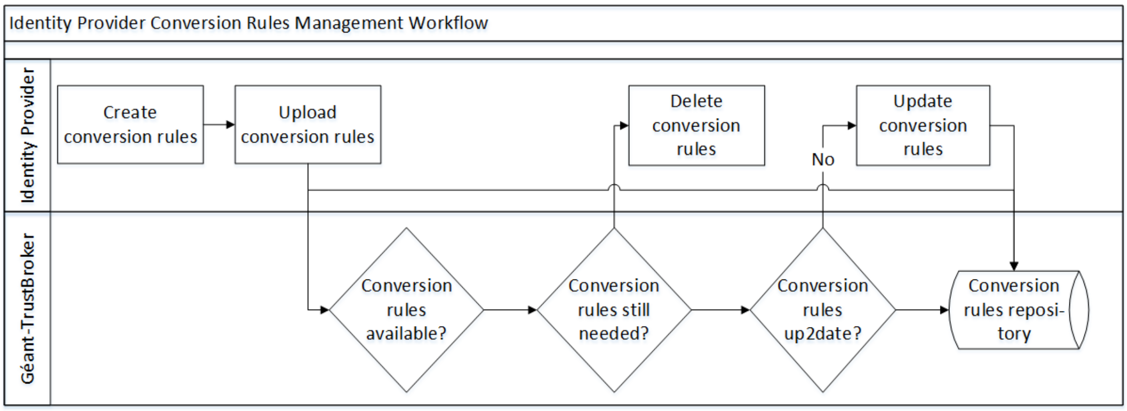
Identity Provider Conversion Rules Management Workflow
Identity Provider
Géant-TrustBroker
Create conversion rules
Upload conversion rules
Delete conversion rules
No
Update conversion rules
Conversion rules available?
Conversion rules still needed?
Conversion rules up2date?
Conversion rules reposi-tory

Figure 4: Conversion Rule Management Workflow

The GNTB API provides the IDP administrator an upload function for his conversion rules with some metadata about the conversion rule, like description, SP federation and name of the conversion rule:

- gntb Conv_Upload(File, Name, SP_ID, Target_List, Description) -> ok, error

The ruleset name serves as an identifier for the conversion rules. The federation of SP parameter is related to the utilized schema to which the rule applies.

Conversion rules are also searchable by the provided API. In order to find a suitable one, the entity-ID of the SP and the federation or schema of the IDP has to be known:

- gntb Conv_SearchForEntity(Entity_ID, Target_List) -> conv, null, error

The entity_ID parameter belongs to the SP for which services the rules apply. Because IDPs, which are members of the same federation uses usually the same schema, only information about the IDPs' federation is needed.

As described above outdated conversion rules have to be checked, possibly updated and activated automatically or after administrator's manual approval. IDP administrators can also retrieve them periodically. This can be set in the IDP's software.

Rules could be identified by its rule_ID or an unique rule name. If an IDP administrator searches for appropriate conversion rules, sometimes more than one ruleset could be available. Thus to support the IDP administrators a rating functionality to assess the quality of rules should be provided, based on different criteria as currency of the rule or how often this specific rule has been re-used by other IDPs.

6. CURRENT STATUS AND OUTLOOK

Géant-TrustBroker's workflows have been described and the the API and data model have been designed. In the current project phase, the second core functionality of GNTB, i.e., the attribute conversion rule repository, is designed, i.e. the storage and exchange of data conversion rule sets that enable on-the-fly data conversion between IDPs and SPs in separate federations. The internal data model defines how conversion rules can be implemented and which allows for the "smart" re-use of shared conversion rules by several IDPs. The API for the conversion rule repository was partly regarded during the design of the overall data model and API.

In the next step the protocols for the communication between the different engaged parties, i.e. SP, IDP and GNTB, as well as the GNTB protocol are specified in detail. A formal specification of the GNTB core workflow will also be prepared as Internet-Draft for the standardization body IETF to enable broad adoption of the GNTB functionality beyond our prototype implementation.

A demonstrator based on Shibboleth with proof-of-concept implementations of

- Géant-TrustBroker service itself,
- the Shibboleth SP GNTB plug-in. and
- the Shibboleth IDP GNTB plug-in

will show the functionalities of GNTB in a local Shibboleth testbed that provides multiple SPs and IDPs in different federations. The experiences gained using this testbed will be used to create a documentation that describes how to set up the Géant-TrustBroker service and the SP/IDP parts as a basis for further deployments in Géant. Currently the piloting of GNTB with several research communities as part of the next overall Géant project phase, starting 2015, is focused in order to provide a long-term working service that complements eduGAIN in the future.

7. ACKNOWLEDGMENT

The research leading to these results has received funding from the European Community's Seventh Framework Programme under grant agreement no 605243 (Multi-gigabit European Research and Education Network and Associated Services – GÉANT).

The authors wish to thank the members of the Munich Network Management (MNM) Team for helpful comments on previous versions of this paper. The MNM-Team, directed by Prof. Dr. Dieter Kranzlmüller and Prof. Dr. Heinz-Gerd Hegering, is a group of researchers at Ludwig-Maximilians-

Universität München, Technische Universität München, the University of the Federal Armed Forces, and the Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities.

8. References

Cantor, S., Kemp, J., Philpott, R., Maler, E. (2005). *Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0*. OASIS Security Services Technical Committee Standard

Géant (2014). *eduGAIN Homepage*. Retrieved January 30, 2014, from: http://www.geant.net/service/eduGAIN/Pages/home.aspx

Hämmerle, L., Schofield, B. (2013). *eduGAIN - Are we there yet?*. Retrieved January 30, 2014, from: https://refeds.org/meetings/oct13/slides/eduGAIN-at-FIM4R-20131002-bas.pptx

Harding, P., Johansson, L., Klingenstein, N. (2008). *Dynamic Security Assertion Markup Language*. IEEE Security & Privacy, no. 2, vol. 6, 83-85

Hommel, W. (2009). *E-Learning in Shibboleth-based federations: The design rationale behind the German DFN-AAI E-Learning Profile*. Proceedings of EUNIS 2009

Solberg, A. (2010). *Dynamic SAML*. Retrieved January 30, 2014 from: https://rnd.feide.no/2010/02/18/dynamic_saml/

SWITCH (2014). *SWITCHaai Resource Registry*. Retrieved in January 30, 2014, from: http://www.switch.ch/de/aai/support/tools/resourceregistry.html

SWITCH (2014). *uApprove – User Consent Module for Shibboleth Identity Providers*. Retrieved in January, 30, 2014, from http://www.switch.ch/de/aai/support/tools/uApprove.html

Terena (2013). *peer 0.11.0: Python Package Index*. Retrieved in January 30, 2014, from: https://pypi.python.org/pypi/peer/0.11.0

9. Authors' Biographies

Daniela Pöhn received a university diploma degree in Computer Science from the University of Hagen, Germany, in 2012.

She was engaged in the IT industry as a full-time software developer during her studies, before she joined LRZ as a Ph.D. candidate in September 2012.

She is involved in the identity management research activity (JRA3 T1+T2) in Géant3+ since April, 2013. The focus is mainly on account linking.

Furthermore she is involved in Géant-TrustBroker.

Stefan Metzger is member of the communication networks planning group at the Leibniz Supercomputing Centre. He holds an international CISSP certification. As head of the LRZ security working group his main focus lays on ISO/IEC 27000-based security management.

He attained in 2005 a Diploma degree in Computer Science from Technical University Munich. Currently he's doing his PhD in security management in large-scaled, inter-organizational infrastructures and offers lab courses in security at Ludwig Maximilians University (LMU) Munich.

Wolfgang Hommel is the Chief Information Security Officer of the Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities, where he is also the head of the communication networks planning group.

He studied computer science at Technische Universität München and has a Ph.D. as well as a postdoctoral lecture qualification from Ludwig-Maximilians-Universität in Munich, Germany, where he teaches information security lectures and labs. His research, for which he was granted the Karl Thiemig foundation's young academics award in 2011, focuses on information security and IT service management in large-scale and inter-organizational scenarios.