Géant-TrustBroker: Dynamic, Scalable Management of SAML-Based Inter-federation Authentication and Authorization Infrastructures

Daniela Pöhn, Stefan Metzger, Wolfgang Hommel

Géant-TrustBroker: Dynamic, scalable management of SAML-based inter-federation authentication and authorization infrastructures

Daniela Pöhn, Stefan Metzger, Wolfgang Hommel

Leibniz Supercomputing Centre, Munich Network Management Team
Boltzmannstr. 1, 85748 Garching n. Munich, Germany
email: [poehn,metzger,hommel]@lrz.de

Abstract. We present the concept and design of Géant-TrustBroker, a new service to facilitate multi-tenant ICT service user authentication and authorization (AuthNZ) management in large-scale eScience infrastructures that is researched and implemented by the pan-European research and education network, Géant. Géant-TrustBroker complements eduGAIN, a successful umbrella inter-federation created on top of national higher education federations in more than 20 countries world-wide. Motivated by experiences with real-world limits of eduGAIN, Géant-TrustBroker's primary goal is to enable a dynamic and highly scalable management of identity federations and inter-federations. Instead of eduGAIN's federation-of-federations approach, Géant-TrustBroker enables the on-demand establishment and life-cycle management of dynamic virtual federations and achieves a high level of automation to reduce the manual workload for the participating organizations, which so far is one of the most significant obstacles for the adoption of Federated Identity Management, e.g., based on the SAML standard. We contrast Géant-TrustBroker with other state-of-the-art approaches, present its workflows and internal mode of operations and give an outlook to how eduGAIN can be used in combination with Géant-TrustBroker to solve current AuthNZ problems in international research projects and communities.

Keywords: Federated Identity Management, SAML, Shibboleth, eduGAIN, Inter-Federation, Trust Management, Géant

1 Introduction

Medium-sized and large organizations, such as universities, typically provide dozens of ICT services to their members, e.g., email, file, web collaboration, and print services as well as services specific for the organization and its business processes, e.g., exam management. Usually, a technical identifier – commonly referred to as username – is assigned to each member and all services can then be used by supplying one's username and some sort of credentials, such as a password. While this procedure is common enough to be considered trivial for

each individual user, the organization-wide management of an arbitrarily large number of users and their permissions for all the services can be challenging. Authorization models such as role-based access control (RBAC) and architectural concepts for centralized Identity & Access Management (I&AM) systems have solved most of the related challenges in theory and are successfully implemented in practice by many organizations, usually based on LDAP servers or relational database management systems that are used as user management backend.

Inter-organizational identity management becomes necessary either when an organization's member shall access external services, for example, because a service such as email has been outsourced to a third party provider, or when members of several organizations shall work together on a common project, such as a research project that involves multiple universities and industry partners. Federated Identity Management (FIM), based on standards such as SAML [3] or lightweight approaches like OpenID, assigns each user to her home organization, called Identity Provider (IDP), and technically ensures that services provided by another organization, referred to as Service Provider (SP), can be accessed by authorized users. The set of all IDPs and SPs that collaborate for a specific reason is commonly referred to as federation, and while federations in many industrial sectors consist of only very few members (only one IDP and one SP is not unusual), many national research and education networks (NRENs) operate large authentication and authorization infrastructures (AAIs), i.e., federations with hundreds of organizational members that include most of the country's higher education institutions and commercial scientific services providers.

While geographic and industrial-sector-specific borders for federations are not imposed by FIM technology itself, they have become a reality due to the historic evolution and growth of FIM use in both industry and higher education institutions: Most sectors and countries run their own federation, resulting in the problem that international and cross-sector collaboration is impeded: Neither a researcher from country A nor an industry partner from country B can access an ICT service operated by a university in country C based on existing AAIs. The only pragmatic solutions are to either create new local user accounts for all project participants, which obviously scales for small projects only, or to set up a new federation specific for the given project or community. Either solution increases the overall complexity for IDP and SP operators as well as for the users, who must either use separate credentials for each service or must be aware of which federations they are members of when accessing an external service.

Inter-FIM is the next evolutionary step and, so far, a young research discipline that still lacks resilient results. Enabling users from one federation to access services in other federations turned out to be a problem with conceptual, technical, and organizational aspects. Most issues stem from two characteristics of today's federations:

1. An organization's membership in a federation usually requires a contract, e.g., either with all other federation members or a central federation operator to ensure that all participants are obliged to certain behavior. For example, IDPs must provide high quality user data to avoid SP misuse based on

fake accounts, and SPs must commit themselves to honor privacy and data protection principles.
2. Federations must be built on common technical grounds, i.e., each member must use the same federation technology (e.g., SAML), and the data format used by all IDPs and SPs in the federation must be harmonized, resulting in the so-called federation schema that defines the syntax and semantics of information provided by IDPs about their users, such as name, email address, and language preferences.

The assumption that a world-wide federation could be built is utopian because no agreement on a common technology, membership criteria, and user data format could ever be achieved for tens of thousands of organizations [8]. A more promising approach is to integrate existing federations into a higher-level umbrella inter-federation: eduGAIN [1] is a successful attempt to set up an inter-federation for NRENs' country-specific AAIs that has been initiated by the pan-European research network Géant. It spans more than 20 federations already, but it grows slowly, supports only a minimalistic data schema, brings additional contractual complexity, and requires significant technical effort for each participating organization.

The limits experienced with eduGAIN in the real world have motivated a complementary and fundamentally different approach to enable Inter-FIM for international research projects and communities in the future. On the one hand, it is much more dynamic and scalable, but on the other hand it cuts back regarding formal contracts while still ensuring a suitable degree of reliability concerning the participant's behavior. While eduGAIN is a federation-of-federations, our new approach, named Géant-TrustBroker, creates dynamic, virtual federations that overcome many organizational and technical issues of other Inter-FIM approaches. In Section 2 we present the concept and goal of the Géant-TrustBroker service and contrast it with the current state of the art. Section 3 then details the Géant-TrustBroker Inter-FIM workflows. The paper is concluded by an outlook to how eduGAIN and Géant-TrustBroker will collude and a summary of the results achieved so far.

2 TrustBroker's distinguishing design and related work

FIM enables SPs to delegate user authentication to each user's IDP and to retrieve certain information about users, the so-called attributes, from this IDP. This workflow implies organizational and technical prerequisites:

- SP and IDP software, operated by different organizations, must be able to communicate with each other:
  - The communication endpoints of both services need to be known and technical information, such as X.509v3 certificates for digital signatures and encryption, must be available. This is commonly referred to as (IDP and SP) metadata.

  - The syntax and semantics of exchanged data must be defined by a common data schema.
- SP and IDP must trust each other: The SP must be able to rely on the data provided by the IDP and the IDP must be confident that the personal data provided about its users is not misused. Traditionally, this requires written contracts, although more dynamic and easier to manage approaches like the Géant Code of Conduct [14] gain importance.

Because the scalability would significantly suffer if each SP had to bilaterally set these prerequisites up with each IDP, federations have successfully become a means to group all SPs and IDPs that share common properties, for example being related to a country's higher education infrastructure. Federations aggregate the metadata of all their participating SPs and IDPs, specify a common data schema, and provide a contractual framework that ensures basic properties for trusting each other.

Merging federations or putting them under the umbrella of an inter-federation is complex in practice due to the heterogeneity of the existing federations: Typically, both the federation data schemas and contract contents differ significantly even if they use the same FIM technology. Forcing all member organizations of a federation to change their data schema does not work in the real world, leading to inter-federation data schemas that are the common denominator of all involved federations, which in turn means that SPs, which require certain user attributes not included in the inter-federation data schema, cannot be used with their full functionality. The additional contracts required between federations and their members make the overall inter-federation more complex and cumbersome to manage. Yet, with major efforts eduGAIN successfully established such an inter-federation. However, in addition to the resulting problems outlined above another real-world problem has not been foreseen: Aggregating the XML-based SAML metadata from many national research federations leads to a huge inter-federation metadata file, whose processing becomes so cumbersome that many of the deployed SP and IDP software packages are slowed down to a crawl that must be either compensated through new hardware investments by all IDPs and SPs or leads to significantly reduced usability for the end users.

Géant-TrustBroker (GNTB) enables the exchange of user identity data across federation borders with the following key characteristics:

1. GNTB provides SP and IDP metadata in an on-demand manner: Instead of distributing the complete aggregated inter-federation metadata to all SPs and IDPs, GNTB provides IDPs only with the metadata of SPs used by at least one of their users and vice versa for SPs. This effectively avoids performance bottlenecks.
2. GNTB enables the exchange of data conversion rules in addition to the other metadata. Instead of supporting only a small set of common user attributes, this allows for the use of arbitrarily complex data schemas, while still ensuring that conversion rules must only be implemented once for each federation and not by each individual SP or IDP.

3. GNTB automates the technical integration of new metadata on the SP as well as on the IDP side when an IDP's user logs into an SP for the first time. This eliminates the manual workload for SP and IDP administrators and avoids waiting times for the end users before they can use a new SP.

The third property means that the technical setup of SP-to-IDP relationships can be fully automated, but it does not have to. Whether full automation is desired or not actually depends on the involved organizations' requirements for trust built on organizational measures:

- On the one hand, SPs of commercial services, which require payments and therefor liability, usually will not accept new users from previously unknown IDPs before they also have a complmentary, mutually signed contract. On the other hand, SPs of free-to-use services, such as a Wiki collaboration web server operated by a university for its research project partners, will prefer easy and quick account rollout for their users and a minimum amount of work to put into user management for the service.
- In practice, most IDP administrators will prefer a fully automated setup because up to now it is a very tedious task that is done anyway whenever one of their users requests access to a new SP. This holds true at least for the higher education sector, where the use of many external services is very common, e.g., due to many inter-organizational research projects. However, those IDP organizations, which are more restrictive about the use of external services, will not want full technical automation, at least not without an explicit manual approval step.

The first of the GNTB characteristics mentioned above regarding the metadata exchange is heavily influenced by related work. For example, most national research federations provide facilities for web-based management of SP and IDP metadata. One advanced example is the Resource Registry (RR) of the Switch federation SWITCHaai [4]: RR provides a web-interface for IDPs and SPs to register their metadata. It allows service providers to specify which of the federation schema attributes they actually use, a seemingly very basic information that is, however, not provided in most other research federations. In return, IDPs have the option to describe all attributes they actually offer [5]. Despite the wide range of provided functionalities, the webtool itself requires manual work for configuration and waiting time for the administrator to receive a basic attribute filter, which he can adapt.

Metadata aggregation and distribution has been designed and implemented by means of the Metadata Distribution Services (MDS) [16] in eduGAIN [1]. The metadata is first aggregated at the federational level, before MDS aggregates and signs the metadata for the whole inter-federation. Entities establish a static bilateral trust relationship, while the Interoperable SAML Profile [18] addresses the exchange of SAML messages. As huge metadata files affect performance and hardly any organization needs the metadata of all other inter-federation members for production – for example, an SP usually never needs information about

all the other SPs in the inter-federation – Dynamic SAML (DSAML, [7]) and Distributed Dynamic SAML [17], developed by Internet2, simplify the discovery of another entity, but does not solve the attribute conversion and attribute filter problems. For the initial trust establishment the metadata consumer validates the signature using a root certificate and establishes the trust. Despite the dynamic character, the metadata has to be published or registered at a central point, e.g. MDS. The Metadata Query Protocol by Young, currently submitted as IETF Draft [15], suggests how to retrieve metadata from entities using simple HTTP GET requests. It therefor solves the problem of huge aggregated metadata files, but otherwise has the same drawbacks as DSAML: attribute conversion, attribute filter and the initial trust establishment require manual work resulting in waiting time for users. The Metadata Query Protocol is one piece of the Metadata Exchange Protocol (MDX). Entities pick a registrar for their metadata and receive attributes from partner entities from one or more aggregators. Aggregators and registrars are linked in order to exchange metadata with each other, analogical to DNS. Similar to MDS, the PEER project [10] implemented a public endpoint entities registry that supports SAML but also other metadata. PEER can obtain metadata from an MDX implementation. Though PEER moves from a huge metadata aggregator to a central system, where administrators can register their domain, many manual steps are needed, e.g. to generate an attribute filter adjusted to the IDP or to establish technical trust between two entities.

3 TrustBroker concept and workflows

GNTB is basically a service to store and retrieve SP and IDP metadata as well as user attribute data conversion rules on demand. The main challenge is to seamlessly integrate the use of this functionality by both SPs and IDPs into standard FIM workflows; to this end, GNTB is tailored for SAML, which is the FIM standard most widely used in research federations, but it could be adapted to other FIM protocols without changing the core functionality.

We distinguish between management workflows and the so-called GNTB core workflow. Management workflows are used by SPs and IDPs to register, update, and delete their own metadata as well as conversion rules in the GNTB registry, similar to how they have to manage their metadata in their research federations. Registering one's metadata is required before the SP or IDP can make use of the GNTB-enhanced workflows: Although self-registration steps initially were and still could be integrated into a single core workflow, this turned out to make it unreasonably complex.

To explain the GNTB core workflow, depicted in Figure 1, let us assume that user Alice from IDP I in federation 1 wants to make use of a service located at SP S in federation 2. As often seen on SP websites, the login / user authentication form at S presents a list of IDPs, which S already knows. However, as I and S have no bilateral technical trust relationship established yet, Alice cannot choose I from that list and instead initiates the core workflow by choosing GNTB as

her IDP. Using standard SAML mechanisms, Alice's web browser is redirected to the GNTB service; similar to federation's SAML IDP discovery services (also known as Where Are You From? service), GNTB then presents a list of the registered IDPs and Alice has to pick the one she wants to use; an account chooser application similar to OpenID's accountchooser.com or equivalent could be integrated as well. Her choice can be remembered, e.g., using cookies, if she always wants to use the same IDP, but account or IDP choosing functionality is increasingly requested by users with several accounts at different IDPs. GNTB passes the information about which IDP has been chosen back to S afterwards.

In the next step, S determines whether an user form IDP I is acceptable or the login should be aborted; for example, the SP could use a blacklist of unwanted IDPs or a whitelist containing only IDPs the SP has contracts with. Also, if Alice has chosen an IDP that is already known to S because she missed it in the list earlier, a regular FIM authentication workflow is started without any involvement of GNTB. If S confirms its interest in users from I, the GNTB core workflow continues as follows:

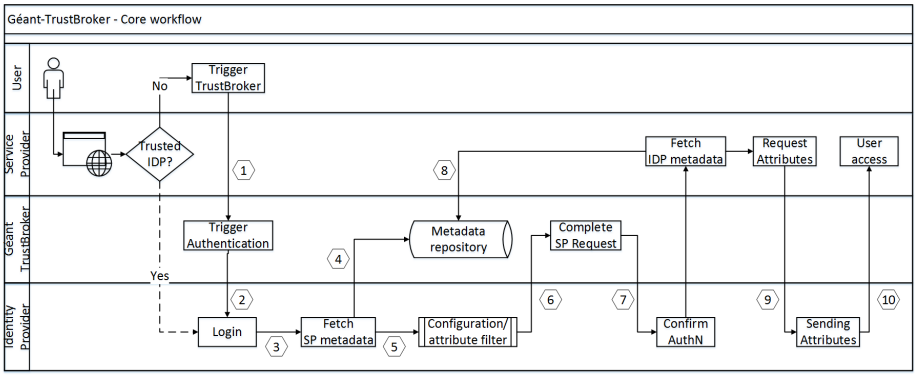

Fig. 1. GNTB's core workflow

1. S prepares a SAML user authentication request, but as it cannot communicate with I directly yet due to missing IDP metadata, it sends the request to GNTB, which temporarily stores it for use in step 7. This is necessary because GNTB must first trigger I to authenticate Alice and determine whether S is an acceptable SP. Otherwise, malicious users could add arbitrary IDPs' metadata to any SP and vice versa, even if they had no valid user account at these IDPs (flood protection).
2. GNTB redirects Alice for authentication to the login page of her IDP I; during this step and step 6, GNTB acts like an SP towards I.
3. In the case of successful authentication or if Alice is already logged in (i.e., a session exists), I fetches S's metadata and attribute conversion rules from GNTB. Otherwise the workflow is aborted showing an error message to Alice.

4. Based on this information retrieved from GNTB, I can automatically update its configuration by adding S's metadata.
5. For the creation of so-called attribute filters, i.e., rules about which user attributes I will send to S on request, I has to check whether it needs attribute conversion rules, which are part of the configuration file attribute-resolver.xml, and whether suitable rulesets are available at GNTB.
6. With Alice successfully authenticated and I completely set up for communication with S now, Alice is redirected back to GNTB, closing the sub-workflow started in step 2.
7. GNTB now redirects Alice back to I again, but uses S's request that was stored in step 1. Since Alice has already been authenticated at I in step 3, I can immediately send a SAML authentication assertion back to S, which also involves redirecting Alice's browser to S again.
8. Because SAML assertions, i.e., the data I sends to S, are usually digitally signed using public key encryption, S now needs to fetch I's metadata, which includes I's public key(s), from GNTB, and add it to its local configuration file of trusted IDPs in order to verify the signature.
9. S now knows that a valid user from I has logged into its service, but it has no other information about the user yet. However, S now has all the metadata required to directly contact I and request SAML attribute assertions that provide some more details about the user.
10. Any other add-ons to the SAML-based user attribute exchange can still be used. For example, IDPs use plugins that ask for the user's permission before sending personal data to an arbitrary SP. GNTB is out of the loop in this stage and does not interfere with existing IDP and SP configuration.

One key aspect here is that the whole workflow is triggered by the user, i.e., the user is enabled to technically connect SPs to her IDP that had no previous interaction with each other. Variants of the workflow explained above exist, e.g., to include manual approval steps. However, if both sides abstain from manual intervention, the user can immediately start to use the service afterwards and does not have to wait on both the SP and the IDP to set up the technical configuration, which we consider a significant improvement over the manual process that is used in all federations so far.

3.1 Variations of the Géant-TrustBroker core workflow

The GNTB core workflow as it has been shown in the previous section covers the primary use case GNTB has been designed for. Additionally, there are several variations of the workflow to handle the following special conditions and constellations:

1. IDP and SP already are members of the same federation.
2. IDP is connected to subordinate Attribute Authorities (AA).
3. SAML Entity Categories are used.
4. SP honors the Géant Code of Conduct.

In this section, we outline the effects of each of these variations. The core workflow is simplified when both the SP and the IDP are already members of the same federation (1) because no user attribute data conversion will be necessary and the required metadata will already be available at both the SP and the IDP. However, the IDP may not yet have been manually configured to send all of the required user attributes to the SP, so at least the IDP administrators could benefit from the automation that can be achieved using the GNTB workflow.

Variation (2) is intended for the growing number of cases in which SPs require information about users which the IDP cannot provide without querying a third party, usually referred to as AA. For example, the use of high performance computing resources via Grid middleware has been FIM-enabled, but Grid SPs typically require user attributes such as Grid user certificate distinguished names (DNs), which most universities do not store in their central I&AM system, so this information cannot be provided by the university's IDP. However, the IDP can retrieve these attributes from an AA; in this case, attribute conversion is necessary if the AA does not use the same data schema as the SP.

Variation (3) applies to SPs that use Entity Categories as described by REFEDS [12] and the Internet-Draft of the IETF Network Working Group [13]. SPs are categorized and IDPs can simplify their setup by applying their configuration to whole categories instead of individual SPs. Therefor, the Entity Category is stored at the GNTB along with the metadata and additional attribute information.

To facilitate the IDP-side trust building process, Géant recommends the use of the Code of Conduct (CoC) [14]. The CoC is a set of privacy and data protection obligations that is closely related to European data protection acts. Its basic idea is that SPs can signal that they honor the CoC and then IDPs can be configured to send personal user data to the SP without the formal requirement of written contracts that govern data protection measures. Variation (4) covers the use case that an IDP wants to check whether the SP honors the CoC, for example, to either reject SPs that do not use CoC or at least require a manual approval step.

3.2 Géant-TrustBroker management workflows

To simplify the GNTB core workflow, SPs and IDPs have to register their metadata and attributes information before the core workflow can be triggered for the first time by any user. Figure 2 shows this workflow for the SP side:

1. The SP has to create its metadata and specify required attributes. The XML-based data format for metadata is standardized by SAML and contains information about the necessary communication endpoints as well as the SP's server certificates, which are used for the verification of XML signatures and encryption purposes by the IDPs. Additionally, the SP can submit a list of required user attributes, which can be marked as either mandatory or optional; if an IDP cannot deliver one of the mandatory attributes, the user will not be able to use the service. A versioning storage backend is used

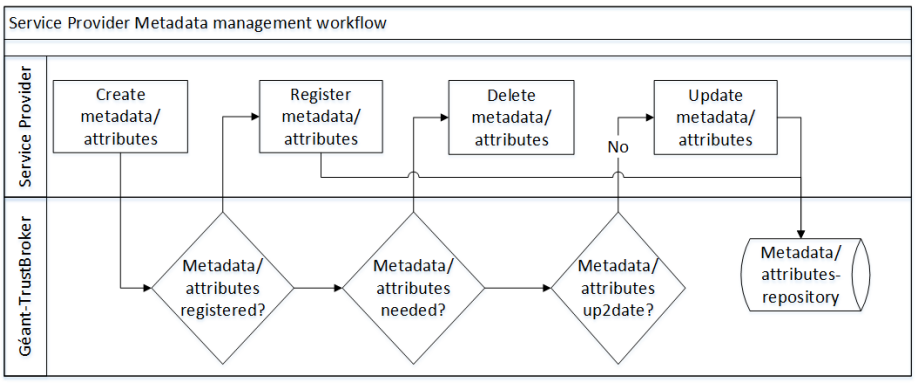

Fig. 2. Géant-TrustBroker management workflow for service provider metadata

in GNTB to track metadata changes because, e.g., certain IDP-side management workflows depend on whether an SP's metadata has changed since it was last retrieved. Write access to GNTB generally requires authentication to ensure that each SP and each IDP can only modify its own metadata. Similar to federation metadata management in most national research federations, technical contacts from SPs and IDPs need to register at GNTB and are assigned an identifier and credentials / keys for the GNTB API.

2. For a first-time setup, the SP uploads its metadata for registration at GNTB. Otherwise, the SP determines whether GNTB has the most recent version of the metadata and can otherwise update it. Metadata changes regularly, e.g., because the X.509v3 certificates contained have a limited validity period and then need to be replaced. Just like updating its metadata in each federation it is a member of, the SP is responsible for keeping its metadata up-to-date at GNTB.
3. If the SP does not want to further use GNTB, it could finally delete its data from the metadata/attribute repository. This cleanup step is optional because the SP must allow GNTB-based login to its service explicitly anyway. Also, GNTB removes outdated metadata entries periodically as well; for example, SP or IDP metadata with X.509v3 certificates whose validity period has expired cannot be used anymore and is purged so that no outdated information can be downloaded by IDPs.

The IDP metadata management workflow is very similar to this workflow and is therefor not described in detail here.

Whenever an SP updates its metadata, the IDPs that have users at this SP need to retrieve the update. Complementary, SPs need updated metadata from those IDPs their users are from. GNTB supports both a push- and a pull-model for transporting updated metadata. Because GNTB stores information about which IDPs have downloaded which SP metadata, they can send the update through a simple GNTB API call. Alternatively, GNTB can be configured as a

standard Metadata Provider service in the IDP configuration, which results in the IDP software automatically polling for metadata updates, e.g., once per hour. Although the latter option causes a delay before the new SP metadata becomes active at the IDP, periodically downloading metadata is the most widely accepted and predominantly deployed method in today's research federations. It gives a higher degree of control that nothing important in the IDP configuration gets overwritten accidentally and is currently preferred by most IDP administrators. This workflow can be seen in Figure 3.

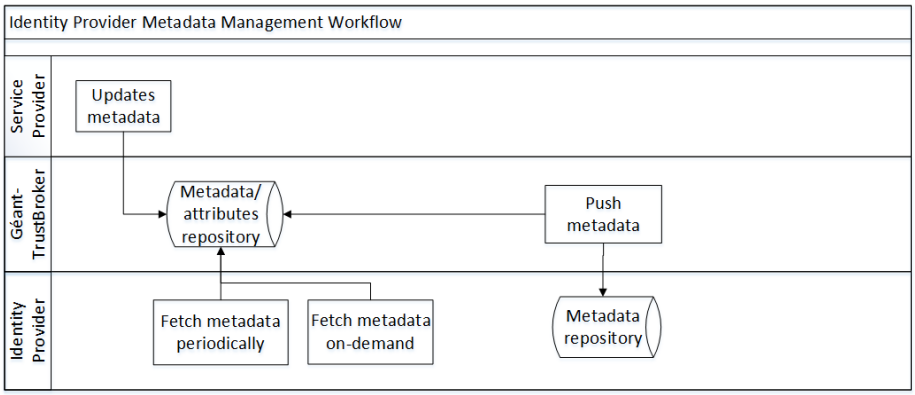

Fig. 3. Géant-TrustBroker workflow for conversion rules

3.3 Géant-TrustBroker workflow for managing conversion rules

In international and cross-sector projects it is not unusual that an SP needs certain user attributes that are not used in the IDP's data schema, i.e., IDPs cannot provide these attributes without additional configuration. However, inter-federation schema discrepancies can often be mitigated using simple data conversion rules. In the simplest case, the attribute only has a different name in both involved schemas, e.g., *surname* vs. *lastname*, which can easily be mapped on a 1:1-basis. The most frequently required rule sets compose a new attribute out of several existing ones – for example, *fullName* is composed from *givenName* and *surname* – or use simple string operations to modify the syntax of existing attributes, e.g., the user's date of birth needs to be converted from *yyyy-mm-dd* to *mm/dd/yyyy* format. Although arbitrarily complex conversion rules could become necessary in theory, these three basic operations – mapping of attribute names, composing new attributes, and string operations for reformatting – are sufficient for almost all Inter-FIM real-world use cases as of today.

GNTB can optionally be used as a central conversion rule set repository by all registered IDPs. The intention is to enable the sharing of implemented conversion rules because usually other IDPs in the same national research federation can

work with exactly the same conversion rules as all IDPs in the same federation are based on the same data schema. It therefor is sufficient if one IDP per federation implements conversion rules for a new SP and makes them available to the other IDPs. The workflow for sharing conversion rules is shown in Figure 4:

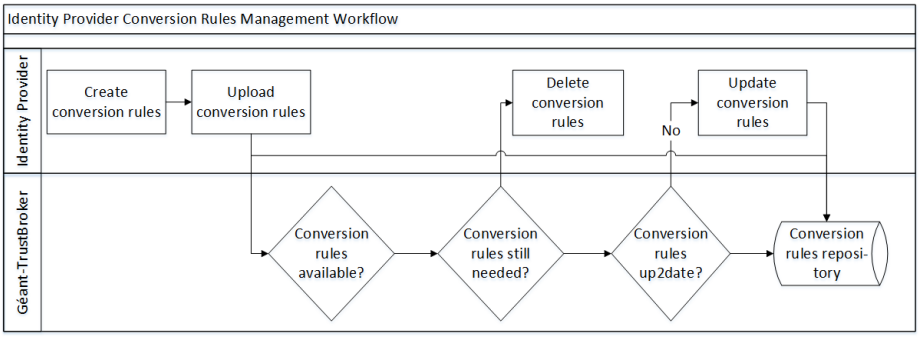

Fig. 4. Géant-TrustBroker workflow for conversion rules

- The conversion rules are implemented and tested by the IDP administrator before the decision to share them is made.
- The IDP administrator uploads the conversion rule set to GNTB. Each rule contains metadata specifying for which SP or SP entity category the rule is intended and which source data schema it is built for. Most national research federation data schemas already have an official, globally unique identifier assigned to them. If no schema identifier is given, GNTB automatically assigns an IDP-specific identifier; this clearly limits automated re-use of the conversion rule, but administrators of other IDPs can still use it if they are confident that the sharing IDP uses the same data schema as they do. Write access to rules is restricted, so IDPs can create, modify, and delete their own conversion rules only for obvious reasons.

Sharing conversion rules on the one hand and re-using some other IDP's conversion rules on the other hand has several implications that must be considered:

- Instead of re-using existing conversion rules, another IDP in the same federation could implement and share conversion rules for the same SP, resulting in multiple, hopefully equivalent, conversion rule sets for the same purpose. During the automated integration of conversion rules at a third IDP, a decision on which rule set to use must be made automatically, defaulting to the newest rule set available. At the moment, GNTB only supports a manual cleanup of duplicates, but other mechanisms – such as a reputation rating system – can be implemented in the future.
- SPs can update their metadata or signal that they need additional or modified attributes. In this case, existing shared conversion rule sets might not

work anymore or not cover the complete set of attributes required by the SP. If this happens, the affected shared rule sets are marked as *outdated* and the IDP administrators who shared them are notified. They have to update their shared rules to remove the *outdated* flag. By default, flagged rule sets are not automatically imported by other IDPs.

GNTB's IDP management workflow also ensures that updates of conversion rules that have automatically been added to the IDP configuration will be downloaded and integrated similar to changes in SP metadata. Because conversion rule sharing is a new and experimental functionality for IDPs, practical experience needs to be gained to determine the long-term feasibility and stability of this approach and which other issues may emerge.

4 Conclusion and outlook

Géant-TrustBroker enables the on-demand, user-triggered exchange of SP and IDP metadata and related configuration data, such as user attribute data conversion rules, across identity federations' borders. It facilitates the fully automated technical setup of FIM-based AuthNZ data exchange and therefor significantly reduces the amount of manual implementation efforts required by both SP and IDP administrators. As a consequence, users can immediately start to use new federation-external services and do not have to wait until the SP and IDP administrators have finished this formerly manual setup process.

It must be kept in mind that a fully automated setup of FIM connections between organizations may not always be desired; especially commercial SPs, which require a written contract with IDPs to ensure, e.g., accountability and reliable payment, are not in the target group of our approach. Instead, the goal, as envisioned in the pan-European research and education network Géant, is to have both eduGAIN and Géant-TrustBroker available as management tools for inter-federations in eScience infrastructures. For SP-IDP-connections that require formal organizational trust building measures such as written contracts, eduGAIN will continue to be first choice. However, eduGAIN will be complemented by Géant-TrustBroker for use by multi-national and cross-sector research projects that want their members to access their distributed services easily and quickly without the previously usual organizational and technical overhead.

The Géant-TrustBroker core workflow, which is an extension of the standard SAML authentication and attribute queries, will be formally specified as an IETF Internet-Draft and submitted for standardization as IETF Request for Comments (RFC). The GNTB prototype and the implementation of the SP- and IDP-sided workflows for the FIM software package Shibboleth will be made available as open source and used for pilot operations in Géant in 2015.

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.

References

1. Géant: eduGAIN Homepage. http://www.geant.net/service/eduGAIN/Pages/home.aspx [Online; 17.01.2014]
2. Hämmerle, L., Schofield, B.: eduGAIN - Are we there yet?. https://refeds.org/meetings/oct13/slides/eduGAIN-at-FIM4R-20131002-bas.pptx [Online; 17.01.2014]
3. Cantor, S., Kemp, J., Philpott, R., Maler, E.: Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0. OASIS Security Services Technical Committee Standard (2005)
4. SWITCH: SWITCHaai Resource Registry. http://www.switch.ch/de/aai/support/tools/resourceregistry.html [Online; 17.01.2014]
5. Hämmerle, L., Lenggenhager, T.: AAI Resource Registry Guide. https://www.switch.ch/aai/docs/AAI-RR-Guide.pdf [Online; 17.01.2014]
6. Solberg, A.: Dynamic SAML. https://rnd.feide.no/2010/02/18/dynamic_saml/ [Online; 17.01.2014]
7. Harding, P., Johansson, L., Klingenstein, N.: Dynamic Security Assertion Markup Language. IEEE Security & Privacy, no. 2, vol. 6, 83-85 (2008)
8. Young, I. A., La Joie, C.: Interfederation and Metadata Exchange: Concepts and Methods. http://iay.org.uk/blog/2009/05/concepts-v1.10.pdf (2009)
9. FIN-CLARIN: The Language Bank of Finland - Language Archive Tools. lat.csc.fi [Online; 17.01.2014]
10. Terena: PEER 0.11.0: Python Package Index. https://pypi.python.org/pypi/peer/0.11.0 [Online; 17.01.2014]
11. Johannson, L.: pyFF Documentation - Federation Feeder 0.9.4 documentation. http://pythonhosted.org/pyFF/index.html [Online; 17.01.2014]
12. REFEDS: Entity Categories R&S. https://refeds.terena.org/index.php/Entity_Categories/R%26S [Online; 17.01.2014]
13. Johansson, L., Young, I. A., Cantor, S.: The Entity Category SAML Entity Metadata Attribute Types - draft-macedir-entity-attribute-00.xml. http://macedir.org/draft-macedir-entity-category-00.html [Online; 17.01.2014]
14. Géant: GÉANT Data Protection Code of Conduct. http://www.geant.net/uri/dataprotection-code-of-conduct/v1/Pages/default.aspx [Online; 17.01.2014]
15. Young, I., Ed.: Metadata Query Protocol - draft-young-md-query-01. http://datatracker.ietf.org/doc/draft-young-md-query/?include_text=1 [Online; 17.01.2014]
16. eduGAIN: mds.edugain.org. http://mds.edugain.org/ [Online; 27.01.2014]
17. Harding, P., Johansson, J., Klingenstein, N.: Dynamic Security Assertion Markup Language: Simplifying Single Sign-On. IEEE Security & Privacy, vol. 6, no. 2, pp. 83-85, March-April 2008, doi:10.1109/MSP.2008.31
18. Solberg, A. et. al: Interoperable SAML 2.0 Web Browser SSO Deployment Profile. http://saml2int.org/profile/current [Online; 27.01.2014]