Options for Joining eduGAIN

Input for the DARIAH/DASISH Authentication and Authorization Infrastructures Workshop in Cologne October 17 - October 18, 2013

Foreword

by Peter Gietz, DAASI International GmbH/DARIAH-DE

The aim of Day Two (Federation for eHumanities and eSocial Science workshop) is to give an overview on the different possibilities of setting up a federated environment for eHumanities and eSocial Science communities in Europe based on an Authentication and Authorization Infrastructure (AAI) and supporting Single Sign-On (SSO). This includes the current activities being made in eduGAIN, TERENA, GÉANT 3 plus and FIM4Research, technical and more important organisational and legal issues. The result of the workshop should be a decision about the solution, these communities should follow.

The current text is meant as input for the discussions.

The largest part of this document is taken from:

> Enabling Users - Options for Joining eduGAIN, Last updated: 25-09-2013,
> Draft of a Deliverable for SA5 Task 5 of GÉANT 3 Plus
> Authors: Lukas Hämmerle (SWITCH), Wolfgang Pempe (DFN),
> Contributors: Marina Vermezovic (AMRES), Thomas Lenggenhager (SWITCH)
>
> The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7 2007–2013) under Grant Agreement No. 605243 (GN3plus).
>
> The Discussions within the DARIAH Pilot for the GÉANT 3 Plus activity Enabling Users had influence on that Document.

The rest of this foreword is a short description of the DARIAH-AAI. Similar infrastructures and issues could have been reported as well for other ESFRI projects in the field of Humanities and Social Sciences like CLARIN and CESDA.

The ESFRI Project DARIAH is establishing a research infrastructure for virtual research environments in the fields of the arts and humanities to enhance and support digitally-enabled research. DARIAH is structured and mainly financed by national projects that provide contributions to a common European infrastructure. The technical infrastructure is supported by a Security Assertion Markup Language (SAML) based AAI and it is planned that all DARIAH services will be protected by SAML Service Provider (SP) and be accessible by scholars via their home organisation account. For this aim an infrastructure has been set up consisting of:

- a central LDAP server that contains homeless accounts as well as a groups the memberships of which provide privileges to specific DARIAH services or data.
- a Shibboleth IdP that acts as authentication authority for the homeless accounts and as an attribute authority for all DARIAH users

- a number of Shibboleth SPs that are configured to do attribute aggregation by sending attribute queries to the IdP after a user has authenticated
- A web-based administration portal that allows for managing group memberships, homeless accounts as well as a hierarchical role system that allows for distributed delegation of management rights within this portal
- A web-based self service interface for providing additional data by all users and for password reset for the homeless users

Since this infrastructure was developed by the German DARIAH project (DARIAH-DE) for now, parts of this infrastructure are integrated into the DFN-AAI. On the longer term it needs to be integrated in a Europe-wide federation, so that all European scholars can use it. There are some other ESFRI projects in the fields of humanities and social sciences (e.g. CLARIN and CESSDA) that have similar requirements and aims. DARIAH wants to act together with these projects to reach the aims.

There are several possible ways forward to reach this, and eduGAIN plays a very important role in all considerations.

Within the DARIAH Pilot for the GÉANT 3 Plus activity "Enabling Users", a number of options for having the Humanities and Social Sciences communities act within a Europe-wide federation had been discussed, based on the following text.

Although three basic options are discussed in the following, only the first two options seem to be reasonable choices for the Humanities and Social Sciences communities.

Enabling Users Options for Joining eduGAIN

by Lukas Hämmerle (SWITCH), Wolfgang Pempe (DFN),

1 Introduction

Within the research communities, the need of federated access to services is seen as an essential success factor, especially in the Social Sciences and Humanities (SSH) sector, where parts of the target group are not highly computer affine and just need a very easy access to web-based electronic research tools. The experiences the research communities made within grid computing showed that X.509 certificate-based infrastructures were a major hindrance for wide community acceptance of research tools. Thus federated Identity Management is seen as the only authentication and authorization technology to be acceptable within the SSH community.

The FP7/ESFRI programs of the EU have led to the efforts of constructing long-term Europe-wide research infrastructures, which need to inter-federate to allow for virtual organizations with members from different countries. The interfederation service eduGAIN ([eduGAIN]) is an answer to such a need.

Multinational research projects usually operate their services in different countries. Many of these services require authentication and authorization and could thus benefit from integration into eduGAIN. Enabling eduGAIN interfederation support for these services requires some know-how and efforts by the service operator. Given that the number of services operated by SSH projects is probably higher than for other research projects and given that the number of services is likely to

increase even more, the question is how research projects can efficiently add their services to eduGAIN.

For research projects the following three options how to add services to eduGAIN were identified:

- Option A: Add services via an existing federation
- Option B: Create an own federation
- Option C: Join via a Hub or Proxy

Each of the above options has its advantages and disadvantages and not all of them are suitable for each research group. This document describes the above options in order to help deciding which of them is best suited for a particular research community or case.

1.1 Federated Identity Management and eduGAIN

An identity federation usually consists of multiple organisations (e.g. universities and research institutes) that agree to use a common infrastructure for authentication and authorisation. eduGAIN is a global interfederation service that interconnects multiple identity federations, both technically and legally. It allows a user from one identity federation to access web-based services in another identity federation. eduGAIN aims at connecting all SAML-based research and education based identity federations world wide. More than half of all known academic identity federations are already connected to eduGAIN as of September 2013, see [eduGAINstatus].

The ultimate goal of eduGAIN is that a researcher from a university X in country A can access a service operated in country B by authenticating with his user account issued to him by his university X. The service then not only learns that this researcher is from university X but it also receives further user information. This can for example include a unique identifier, name, email address and other data, depending on what information is requested by the service and what information is released by university X. The identity information (especially the unique identifier) can then also be used to perform authorisation. Authorisation can rely on identity data like the researcher's organisation or affiliation but it is more likely to rely on data managed by a research project itself.

A service, like a web document storage application or a research database, in the context of SAML is protected by a (SAML) Service Provider which enforces authentication and implements the authorisation. In the context of SAML-based federations, all Service Provider of a particular federation have to be listed in that federation's metadata (XML) file. Joining a federation means accepting the federation's policies and agreements. Technically, it means registering the Service Provider with that federation's operator in order to get the SP's description included in the federation's metadata file. The same of course applies to Identity Providers. Identity Providers are those entities that authenticate users of a particular organisation. They usually are connected to an organisation's user directory.

1.2 (deleted)

1.3 Current Issues

Some SSH research groups like CLARIN and CESSDA have already conducted pilot projects where some of their services were added to eduGAIN. They discovered many issues and problems in the beginning. Part of the difficulties probably also were related to the fact that these pilots were conducted at an early stage of eduGAIN. The experienced issues that are relevant for the context of this document are described below.

1.3.1 Lack of Identity Providers

When an identity federation joins eduGAIN this does not mean in all cases that all of its Identity Providers (IdP) and Service Provider (SP) also immediately participate in eduGAIN. In fact, most federations have implemented an opt-in model that leaves each IdP and SP the choice whether it wants to become interfederated or not. The opt-in model was implemented mostly because it makes no sense for all IdPs and SPs of a federation to be part of eduGAIN as they might only be used within a single organisation or a single federation. Also, often some additional efforts are required to enable an IdP or SP for eduGAIN. These efforts may include policy, configuration and application adaptations. Therefore, this takes some time. The opt-in process guarantees that only those organisations are exposed to eduGAIN which have completed these steps.

From the research projects' perspective it is mostly the Identity Providers (IdPs) that are of interest to their services because the IdPs allow the research project participants to use their organisation's identity to get access to the research project's services. However, most projects have participants whose organisations have not yet started or completed the opt-in process. Therefore, their users are not yet able to use eduGAIN to authenticate which forces the eResearch projects to provide alternatives in form of own Identity Providers for "homeless" users, the so-called "Homeless Identity Provider".

It is expected that with time, this issue becomes less of a problem as more and more Identity Providers opt-in - a process which, as a kind of chicken-and-egg situation, will be brought forward by an increasing number of attractive services available via eduGAIN.

1.3.2 Lack of Attributes

Another problem for services making use of eduGAIN is that they often request more attributes than the Identity Providers (e.g. operated by the universities) are willing or able to release about their users. This then causes problems when users want to access a service and authentication fails due missing attributes. Problem most certainly arise if the eduPersonPrincipalName (ePPN), the eduPersonTargetedID (ePTID) and/or email address are not released as these attributes are used to identify users. Collaboration in research contexts / infrastructures requires reliably identified users. That's why for many services anonymous usage is not an option.

There are technical and legal aspects why Identity Providers sometimes don't want to release all the requested information in form of attributes. Generally, the technical problems concerning attribute release can be addressed with better tools (e.g. to get the user's consent for attribute release after login), less privacy/security sensitive attributes (eduPersonTargetedID or the new eduPersonUniqueID) and documentation. To address the legal and policy issues that may hinder attribute release on organisation level, the GÉANT Data protection Code of Conduct (CoC, see [CoC]) was created. Basically it is a declaration stating that the operator of a Service Provider obeys a couple of basic data privacy principles in compliance with the current EU data protection directive. Based on the CoC, Identity Providers then can easier create attribute release rules without risking legal issues when releasing data about their users.

1.3.3 Lack of Level of Assurances

The research projects often would like to know better how the identity vetting of an eduGAIN users was performed and on what basis. However, there is currently no agreed-on attribute that could be used to reliably express such information. Also, it is likely to take years for organisations to harmonize their identity-vetting processes and to provide an agreed-on attribute for all of their users.

Therefore, solutions for this issue can only be implemented on the level of a particular service. This implies that currently level of assurances and improved identity vetting have to be performed by the service itself, which of course is not ideal. Potential solutions would outsource this work to another service (authentication-as-a-service). SURFnet, the research and education network in the Netherlands, plans introducing such a service that allows users to go through a standardized identity vetting process and authenticating with two-factor methods.

If a Service Providers receives user attributes not only from a user's Identity Provider but also from a third-party attribute authority because it requires service/community-specific roles (entitlements), the registration process with the attribute authority provides an additional opportunity to confirm the user's identity.

1.3.4 Lack of support for non-browser applications

Most SAML-based identity federations today support only the web SSO profiles, which limits the application to web-based applications. Some research projects however also have use-cases for non-web applications like SSH access. SAML includes profiles like the SAML Enhanced Client or Proxy Profile (SAML ECP) which could serve as a solution. This solution also would work in eduGAIN. However, this profile has not been deployed yet by many Identity Providers. Therefore, it is hardly usable by the research infrastructure. There are workarounds to get access to non-web resources that involve an initial web authentication. However, they often are not user-friendly and not easy to maintain.

2. Different Options to Join eduGAIN

The following chapters explain the different options how Service Provider (SPs) and IdP(s) can be added to eduGAIN. It is assumed that research projects are mostly interested to reuse the identities administrated by the universities and research institutes for which their project participants are working. Therefore, in the context of research communities, the focus for these research projects is to primarily add SPs, which protect the actual services, to eduGAIN. Most probably, large research projects will however also want to operate at least one IdP that contains identities for users that are not affiliated with a university or a research institute (yet) participating in eduGAIN.

Option A – Add services via an Existing Federation

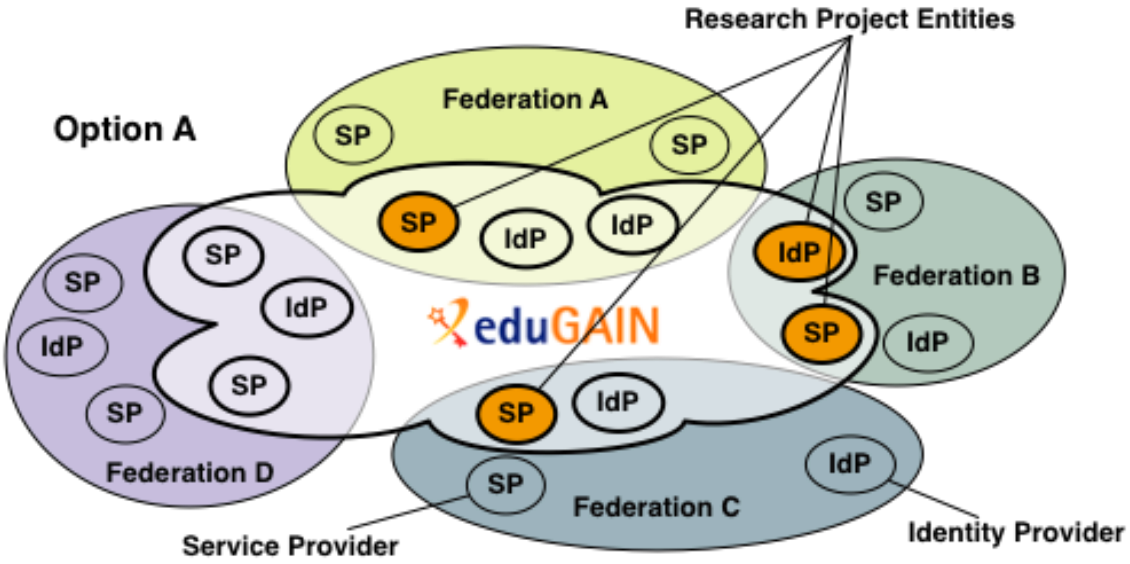

Adding all SPs to eduGAIN via one or more existing federation is the most straightforward option. The registration procedures to add entities to an existing federation vary from federation to federation. The same applies to the steps necessary to enable a service for eduGAIN. In case of a research project with services operated in multiple countries, the registration of the entities probably would be done in the respective federations.

Advantages	Comments / Discussion
Reuse of know-how, infrastructure, documentation and guides, policies, legal framework, processes of existing identity federations.	When the services of a research project are registered in the country they are operated in, the administrators registering the services might already be familiar with the processes of registering SPs. They will get support and assistance from their local federation operators in the language they speak. Most federation operators are also likely to have a large know-how in the area of SAML and federated identity management in general because many of them have been operating federations for years. Federations operated by National Research and Education Networks (NREN) are also very likely to persist for a longer time as the identity federations have become very important for the NRENs when it comes to offer new services to their community.
Identities not tied to project services only	Services registered with the respective home federations are also available for users from local IdPs which for some reason are not willing or able to interfederate.
From the user's point of view a very	Accessing a service in eduGAIN is no different than

Advantages	Comments / Discussion
transparent solution	accessing a service within the local federation.
Technically straight-forward	The procedure is no different than for any other service in a particular federation that wants to become eduGAIN-enabled.

Disadvantages	Comments / Discussion
Potentially many different federations/processes to deal with when registering SP/IdP, cf. also [REFEDSbarr]	This is primarily an issue if a central service unit of the research project carries out the registration and deployment of all SPs. Also, it is only relevant if different organisations in different countries take responsibility for these SPs even though a central service unit manages them. If the research project operators at their respective home organisations register the SPs and/or Homeless IdPs and if their home organisation is responsible for the services, this should be no issue because the operators of these services only deal with a single federation. Most federations primarily care about who is (legally) responsible for a SP/service. Therefore, an alternative approach could also be that one single organisation could take all responsibility for all SPs by a particular research group, regardless in which country they are operated. This then most likely would allow this organisation (or a service unit acting in the name of that organisation) to register all SPs in a single federation, regardless of whether the SPs are operated. This then would lead to consistent registration procedures as well.
Integration of the Homeless IdPs is still needed, and thus also possibly one user multiple identities	Except for very small research groups it is always the case that the research project includes people that don't have an identity at an organisation that already participates in eduGAIN. Therefore, it is likely that most research groups will have to operate a home for the homeless users. Ideally this

Disadvantages	Comments / Discussion
	homeless IdP then also joins eduGAIN.

Examples

- There are already quite a few research applications part of eduGAIN. Among them science gateways from the African Grid community, INDICATE E-Culture, agINFRA, DECIDE, EarthServer, EUMEDGRID, GISELA and IGI.

Option B – Create Own Federation

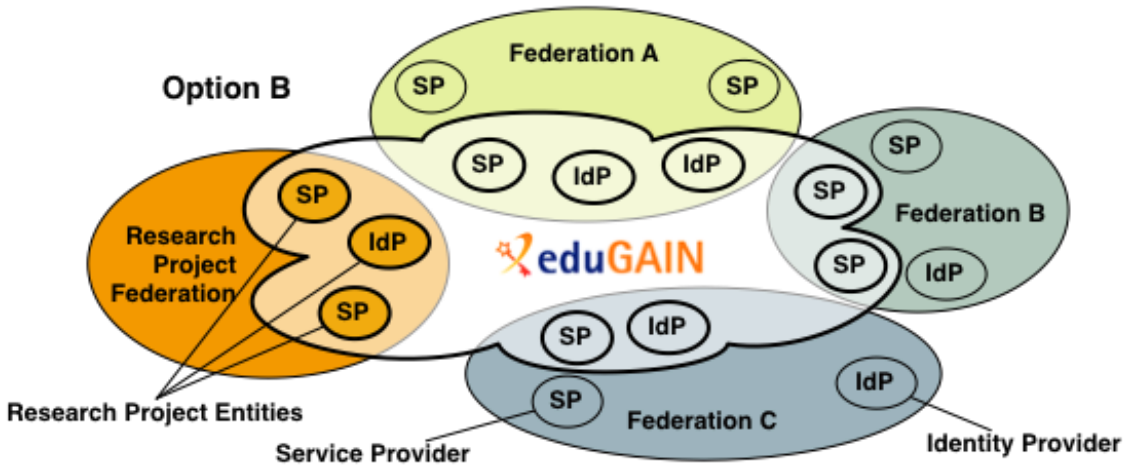

A research project with many services may choose to create an own federation with all SPs and Homeless IdPs of all participating (national) partners and join eduGAIN as a whole federation. Each of the project's entities is registered with this federation according to a uniform set of rules.

Operating a Federation

Operating an identity federation normally involves technical, legal and policy aspects. Most existing identity federations operate at minimum the following services:

- A SP and IdP registration service
- A central Discovery Service used as fall-back in case services don't integrate an own Discovery Service.
- Metadata aggregation tool to process, publish and consume federation as well as eduGAIN metadata
- A home for the homeless Identity Provider for users not (yet) affiliated to a federated organisation

On the legal and policy side of things, most production federations have a legal framework that defines the rights, liabilities and duties of participating Service and Identity Providers. For eduGAIN it is also needed to have a metadata registration practice statement, which describes how entities are registered.

Federations also have to offer various documentation to their community. This includes at minimum how to deploy and configure a Service Provider and how to register it with the federation. As federated identity management often is non-trivial, all federations also operate a help and support desk that assists in case of technical problems. Depending on the size of a federation these helpdesks requires a considerable amount of manpower.

Advantages	Comments / Discussion
Potentially greater influence on IdPs to release attributes if SSH entities join up to create a federation	It is not very likely that IdP administrators notice the registration authority of entities. Who registered a particular entity is well visible in eduGAIN metadata but it is mostly irrelevant to IdP administrators.
Consistent registration of SPs within only one single federation	By creating an own federation and adding all SPs operated in multiple countries to that federation according to own instructions and deployment guides makes the installation and configuration more consistent across the research project. Also see comment on option A.
Representative in eSG → Influence on eduGAIN operations	By creating an own federation and getting accepted as eduGAIN member federation, the research project is granted representation in the eduGAIN Steering Group (eSG), which controls the operation of eduGAIN and accepts news member federation. As of now, each federation has the right to assign one representative.
From the user's point of view a very transparent solution	The user gets displayed information about the service he is accessing and only those attributes are released, which are requested by the service.
Technically straight-forward	Federations have already deployment guides for this scenario as it is no different than registering basically any service within a particular federation.

Disadvantages	Comments / Discussion
Overhead to manage a federation (policies, metadata management, own deployment guides, etc. This requires a sustainable *operating un*it and more or less permanent *legal advice*.	Operating an own identity federation comes at certain costs and requires some persistence. Because research projects typically last only a few years, the overhead for creating and decommissioning a federation seems rather high.

Disadvantages	Comments / Discussion
Integration of the Homeless IdPs is still needed, and thus also possibly one user multiple identities	Except for very small research groups it is always the case that the research project includes people that don't have an identity at an organisation that already participates in eduGAIN. Therefore, it is likely that most research groups will have to operate a home for the homeless users. Ideally this homeless IdP then also joins eduGAIN.

Examples

- None so far. The current SAML-based identity federations are all operated by National Research and Education Networks (NREN). The CLARIN SP federation could be seen as an attempt in this direction though

Option C – Join via a Proxy

Operating a proxy that allows eduGAIN users to access services of a research project might have the advantage that only one Service Provider has to join eduGAIN because all services can be hidden behind the proxy. There are sub-options how to implement a proxy.

Option C.1 – (SAML) IdP Proxy/Hub

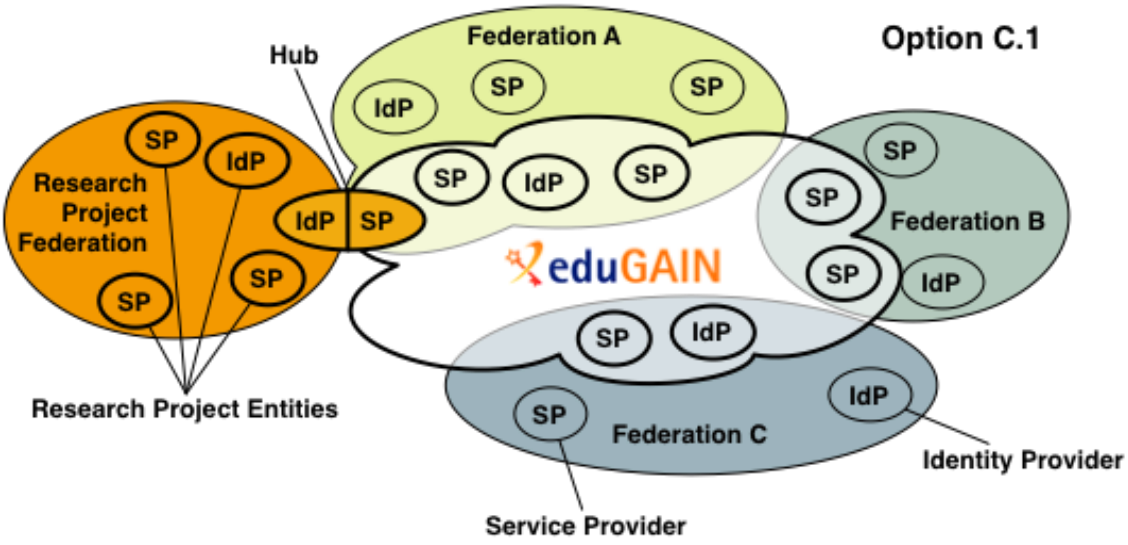

This approach is described in detail on the page:

https://spaces.internet2.edu/display/GS/SAMLIdPProxy

The goal is to build the proxy as a hub that transforms eduGAIN SAML2 assertions to SAML2 assertions used within the research project. The hub consists of an SP facing eduGAIN, and an IdP facing the research project SPs. Optionally, a user directory on the hub is used to store, transform and extend user data. The IdP's SSO login handler would have to be protected by the Service Provider. The Service Provider would have to offer a Discovery Service that could also include any number of Identity Providers (for the homeless users) operated by the research project.

In this scenario one single SP has to be registered in an existing federation that is an eduGAIN member. It might also be necessary to register an IdP if the users managed via the hub also should be able to access other eduGAIN resources not managed by this research group.

Advantages	Comments / Discussion
User data can be extended, transformed, augmented.	As all assertions containing data flow through the hub, it is relatively easy for the hub to modify this data. This could be useful to introduce project internal level of level of assurance or to add group/affiliation attributes that then can be used by the services behind the proxy.
Bridging communities becomes easier	The hub could be extended to support multiple protocols and authentication mechanisms that would allow bridging different research communities and infrastructures (e.g. SAML federations and X.509-based grid community). The advantage would be that such changes have to be implemented only at the hub itself but maybe not at the services behind it.
eduPersonTargetedID [ePTID] would be sufficient for user mapping	The Hub would need at minimum a unique identifier like the eduPersonTargetedID for a user. Getting this attribute should in general be unproblematic, as it does not convey more information about the user than a random string and the Identity Provider where the user authenticated. The user then could add additional attributes himself on the proxy or the proxy could add attributes based on the user's affiliations within the research project.

Disadvantages	Comments / Discussion
Requires development work to implement the bridging/proxying.	The effort to develop and maintain such a proxy should not be underestimated. Currently, there is no out-of-the-box solution that implements proxying easily. The services behind the proxy have to

Disadvantages	Comments / Discussion
	interoperate with the proxy itself, which technically also requires a mini-federation with a minimum set of agreements and policies. Of course these services also need to be specifically configured to interoperate with the proxy.
Hub is a single point of failure.	As all assertions from eduGAIN to the services have to go through the proxy, the proxy becomes a single point of failure. Therefore, it should be ensured that the proxy itself is part of a high availability solution that minimizes downtime.
The proxy hides all services behind it. Because they often have different attribute requirements, the proxy itself has to request the superset of all attributes required.	When users access a service behind the proxy, the Identity Provider at which they authenticate only knows the proxy but not the service behind it. Therefore, the user does not get information on the actual service he is accessing. Requesting more attributes than actually needed is problematic from a data protection point of view.

Examples

- The large photon and neutron research community (CRISP and PanData projects) intend to build a hub like described above in form of the Umbrella infrastructure.

Option C.2 – (Web) Proxy

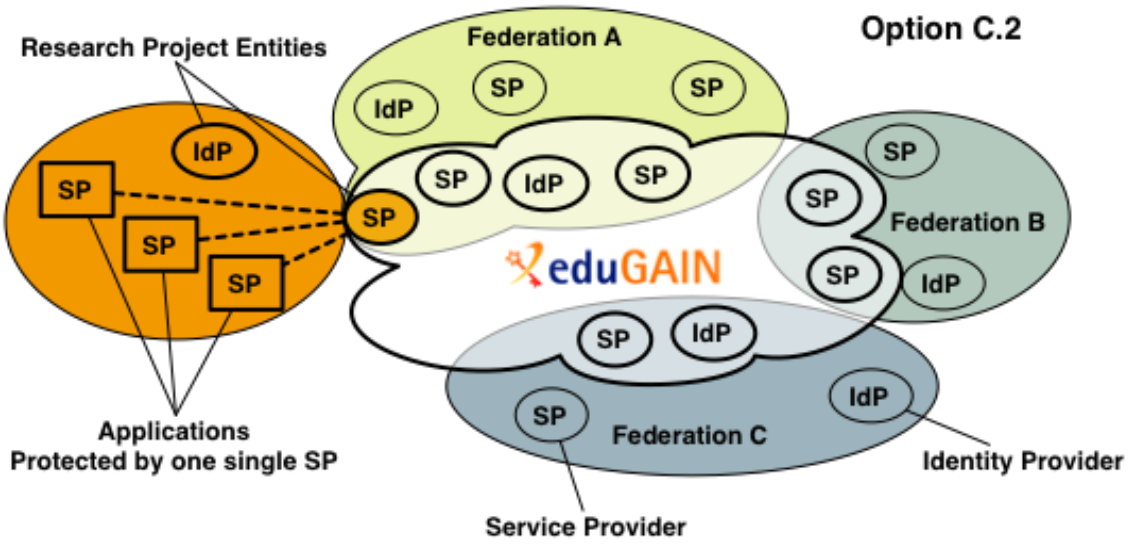

An alternative to operating a full-fletched proxy consisting of an IdP/SP component would be to operate a standard Apache web server with one physically installed SP that is configured for multiple

virtual SPs. Together they could serve as reverse proxy that can protect multiple applications behind it. It must be ensured that all traffic to the applications flows via the proxy and that applications behind the proxy don't accept HTTP headers containing Shibboleth attributes from hosts other than the proxy. The proxy can host multiple Service Providers and they can be registered to one or many different federations. One SP registration could be used to serve multiple applications. But this is intransparent for the user because for example the applications behind the proxy might have different needs for attributes, which then conflicts with data protection principles.

The general concept of this approach is described in detail for Shibboleth on the page:
https://wiki.shibboleth.net/confluence/display/SHIB2/SPReverseProxy

This approach can also be part of the Options A and B.

Advantages	Comments / Discussion
Only one Service Provider would have to be operated to protect multiple applications.	As all services are behind the proxy, it is essentially sufficient to operate one single physical Service Provider that is configured to host multiple virtual/logical Service Provider for the different applications it protects. Operating only one single Service Provider might also have the advantage that fewer people need to be familiar SAML and federated identity management in general. To improve availability it might however be advisable to operate multiple redundant Service Providers.
It might be necessary to register the SP only in one single federation	As the proxy is operated only in one single country, it probably will have to be registered only in a single federation even though the actual services behind the proxy might be operated in various countries.
SP can still be configured for different applications with different attribute needs. Therefore, the proxy becomes transparent for the user.	Depending on the SAML implementation, each virtual/logical Service Provider can be configured individually for each service it protects. Therefore, there are no drawbacks regarding data protection and transparency from a user's point of view.

Disadvantages	Comments / Discussion
Proxy becomes single point of failure.	As all traffic to the services flow through the proxy, it should be made redundant (multiple instances, shared database for session management). Otherwise the outage of the proxy will cause service disruptions for all services behind the proxy.
All network traffic would have to flow through	While this generally is not an issue on a technical level, it of course increases the risk that network

Disadvantages	Comments / Discussion
proxy	problems affect the operation of the service.
Increased complexity and harder to debug	The increased complexity with services distributed on different web servers (behind the proxy) managed by different administrators makes it difficult to well maintain this solution. Also, in case of problems debugging might become harder.

Examples

- There are many universities in various federations that use this approach within a single federation. The difference between this scenario and an interfederation scenario should however be relatively small, as the basic principles stay the same.

Overview Table

Find below an overview table that lists some of the above points:

	Option A Add services via an Existing Federation	Option B Create Own Federation	Option C1 (SAML) IdP Proxy/Hub	Option C2 (Web) Proxy
Technical Overhead	Reuse of infrastructure, documentation, guides and processes of existing federations	Own metadata management (register entities, create/sign/publish metadata file) must be deployed and maintained.	Must implement and maintain the bridging/proxying. Own deployment instructions and metadata managements is needed for the SPs behind the bridge.	Web Proxy (specially configured Apache and SP) must be deployed. Applications behind the proxy must be protected such that they accept only connections from the proxy. Otherwise, HTTP Header spoofing becomes easy.
Administra-tive Overhead	None	Setting up a federation might include creating own agreements, policies, own deployment guides, metadata registration statements etc. In order to join eduGAIN, at least a federation policy	Similar to option B but can be less formal as the proxy would join eduGAIN via an existing federation.	None

		and a metadata practice statement must be available. Federation must be accepted by eduGAIN Steering Group.		
Entity Registration Overhead	Each SP needs to be registered. If there is a large number of SP, potentially they have to be registered in different federations which means many processes to deal with. If the research project operators at their respective home organisations register the SPs and/or Homeless IdPs and if their home organisation is responsible for the services, this should be no issue because the operators of these services only deal with a single federation.	Consistent registration of SPs within only one federation, regardless of where the service is operated. But the SP has to be somehow registered.	At minimum one SP needs to be registered within only a single federation. If identities of that community shall also be able to access other eduGAIN services, an IdP also must be registered.	In a basic scenario one SP needs to be registered. If multiple SPs are used in order to reflect different attribute requirements by applications, they must be registered separately, potentially in one or multiple federations.
Overhead to register own IdP in eduGAIN	Registration with an existing federation depends on the federation's membership rules and might qualify to paying fees.	Easy to register an IdP with an own federation. Criteria according to which IdPs can join federation are created by federation operator.	Registration with an existing federation depends on federation's membership rules and might qualify to paying fees.	Registration with an existing federation depends on federation's membership rules and might qualify to paying fees.

Maintenance overhead	No additional maintenance besides that of operating the individual SPs and IdPs	Requires additional maintenance to manage and operate a federation. Support has to be provided to new SPs and IdPs. They have to be registered. Conformance to policies has to be checked, etc.	Requires additional maintenance of the Proxy/Hub code and configuration. Hub must somehow manage which SPs and IdPs behind the hub are accepted by this hub. SPs and IdPs behind the hub must be supported.	Requires maintenance of the Proxy. Which is basically an Apache web server plus a Shibboleth SP. Both will have a non-trivial configuration and setup.
Resilience to failure	Only depends on application. No single point of failure usually.	Only depends on application. No single point of failure usually.	Proxy is a single point of failure. Should be operated redundantly and with high-availability setup ideally.	Proxy is a single point of failure. Should be operated redundantly and with high-availability setup ideally.
Transparency and Data protection from user's point of view	Good. Normally, no big difference between accessing a service in local federation or via eduGAIN.	Good. Normally, no big difference between accessing a service in local federation or via eduGAIN. If users also get an account for accessing services of the particular research community, this might be confusing as users might use two accounts.	Not very transparent and ideal from a data privacy point of view because the Hub's SP must always request the maximum set of attributes that are behind the proxy.	Good if each application with different sets of attributes is registered individually even though they are protected by the same proxy. It would however also be possible to register multiple applications with the same attribute requirements together as one logical SP. In this case, the transparency would suffer.
Other Aspects		Greater influence on the operation of eduGAIN because all participating federations have a representative in the eduGAIN Steering Group. Might not be suited for individual research projects that last only a few years.	Allows transforming, extending and augmenting user attributes before sending them to services behind the proxy. The proxy could also serve as protocol translator, which would allow to bridge other communities by supporting other protocols than SAML.	Compared to the other solutions only one physical installation of an SP is operated. This approach can also be mixed with Option A and it can be a part of Option B as well as Option C1.

3 Recommendations (deleted)

Glossary

ePPN	eduPersonPrincipalName, cf. [eduPerson] attribute schema
ePTID	eduPersonTargetedID, cf. [eduPerson] attribute schema
ESFRI	European Strategy Forum on Research Infrastructures, [ESFRI]
Federation	Identity federation. An association of organisations that come together to exchange information as appropriate about their users and resources to enable collaborations and transactions.
FP7	Seventh Framework Programme, [FP7]
IdP	Identity Provider. A server acting in an Identity Provider role as defined in SAML 2.0 specifications, cf. [SAMLOverview]
SAML	Security Assertion Markup Language, http://www.oasis-open.org/committees/security
SP	Service Provider. A server acting in a Service Provider role as defined in SAML 2.0 specifications, cf. [SAMLOverview]

References

[CoC]https://refeds.terena.org/index.php/Data_protection_coc

[eduGAIN] http://www.edugain.org

[eduGAINconstitution] http://www.geant.net/service/eduGAIN/resources/Documents/GN3-10-326%20eduGAIN_constitution%20v2.0.pdf

[eduGAINjoining] http://www.edugain.org/technical/joining_checklist.php

[eduGAINmdprofile] http://www.geant.net/service/eduGAIN/resources/Documents/eduGAIN_metadata_profile_v3.doc (final draft)

[eduGAINstatus] http://www.edugain.org/technical/status.php

[eduPerson] eduPerson(200806), http://middleware.internet2.edu/eduperson/

[ESFRI] http://ec.europa.eu/research/esfri/

[FP7] http://cordis.europa.eu/fp7/capacities/research-infrastructures_en.html

[REFEDSbarr] https://refeds.terena.org/index.php/Barriers_for_Service_Providers

[SAMLOverview] https://www.oasis-open.org/committees/download.php/13525/sstc-saml-exec-overview-2.0-cd-01-2col.pdf