Authorisation/Group Management for E-Infrastructures

May 2015

Foreword

A distinctive feature of e-infrastructures is that most individuals' authorisation to access a particular service does not come from their home organisation (as it does for site-licensed journals, for example) nor from the operator of the service (as in traditional, non-federated, access).

Instead, authorisation is largely devolved by service owners to individuals who act as 'group leader' or 'principal investigator' when deciding who else can share their access to a particular service, dataset or experiment. Often the group and its resources may form a virtual organisation, crossing the boundaries of the real-world organisations that employ individuals and operate services. The interface through which leaders create and manage their groups is therefore a key component, effectively defining the membership of and roles within the virtual organisation. This involves interactions with human users as well as networked systems providing both authentication and research services.

This paper considers the various functions and interfaces that might be required of a general group management platform and how they are provided in current e-infrastructures. Most UK and international e-infrastructures currently use dedicated group management platforms, bound to particular infrastructures or user communities, which implement the particular functions and interfaces required by those communities. However, future cross-infrastructure and cross-community research is likely to require group management platforms to inter-operate and provide a wider range of functions. The paper suggests ways that these developments towards a more general service might be facilitated, either by enhancements to individual platforms or, where they require corresponding changes to a number of different infrastructure components, by studies, pilots or recommendations towards a common development roadmap.

"Authorisation/Group Management for E-Infrastructures"

© Jisc
Published under the CC BY 4.0 licence
creativecommons.org/licenses/by/4.0/

Contents

Introduction

For most traditional on-line resources, permission to access the resource ("authorisation") is either granted or denied by the resource owner. If an authentication step is needed to identify those individuals who do have permission – for example, because they have paid a subscription and received a username and password in return – then the resource owner provides the authentication process as well. One of the early applications of federated access management has been for resources provided under site licences: here the licence-holder (such as a university or college), not the resource owner, decides which individuals are authorised to benefit from the licence; typically, as the users' home organisation, it also provides individuals with authentication credentials to prove to on-line services that they are the people to whom authorisation has been granted.

E-infrastructures often use a third scheme. While the authentication process that connects an individual to their on-line account(s) may be provided by either the resource owner or their home organisation (in federated schemes the latter is more common), responsibility for authorising individuals to access particular resources is devolved by the resource owner to an external individual. In research terminology this is often the Principal Investigator (PI) of a project: more generally the function may be referred to as a group manager. Here the terms will be used as synonyms. This pattern of devolved authorisation covers a wide range of different types of resource: for example a research grant of time on a telescope or other equipment may be made to a named PI in the expectation that they will allocate the time appropriately among their collaborators; a researcher may use a research infrastructure to store their raw data and wish to manage access permissions granted to other individuals working on it; space on a collaboration tool such as a wiki or video conferencing platform may be set up by one individual who then wants to enlist others to participate and contribute.

A model that introduces a third party group manager into the access management process requires a new interface for this group manager to enter their instructions.

Although this interface could be limited to managing group memberships and permissions, this paper suggests that there are other related functions for which the group manager is also the authority and that these could naturally and efficiently be provided through the same platform. Through such a group management platform, a Principal Investigator could manage all the stages of a collaboration: inviting participants and determining how they can be authenticated; providing them with access to the necessary resources; supporting accounting and, if necessary, accountability for use and misuse of those resources.

One function that cannot in general be provided by a purely on-line platform is identifying the specific real world individuals that the manager wishes to join the group. That can only be done using knowledge obtained by the manager in the real world – most often what the desired collaborator's e-mail address is. Relying on information obtained within an on-line environment raises issues of mistaken identity (names are not unique), forgery (anyone can set up an account called "Albert Einstein"), and privacy (home organisations may well be reluctant to pro-actively release personal information about their account holders). In fact this apparent limitation should make the process of establishing a federated group management system significantly easier, because it means there is no need for home organisations to release information such as real names that might raise privacy or legal concerns. All that is required from the home organisation's systems is an authentication process that generates a persistent identifier (which may be opaque) for the same user. As discussed below this is sufficient for the user to either accept a personal invitation to join a group or to request such an invitation from the group manager. All additional information can be obtained from either the individual or the group manager who is, after all, the ultimate authority for who the intended collaborators are.

Functions

Considering the group management platform more generally as a "Principal Investigator interface" suggests a group of functions that are likely to be used together and could usefully be provided by a single platform.

In addition to the obvious Group Management, these are Invitation, Alternative Authentication Sources, Shared Management, Accounting and Accountability. These are discussed in the following sections, with a final section showing how the Group Management Platform fits into the overall e-infrastructure.

Group Management

The basic function of a group management platform is to allow Principal Investigators (PIs) to manage groups. In its simplest form this involves managing a list of people who are authorised to access the resources that have been granted to the PI. The service that controls access to those resources will be expected to implement the list, allowing access by those who are on the list and denying access by others. Since group membership is likely to change in time, the PI should be able to add individuals to the list and remove others from it. Web interfaces are often used to provide the PI with a convenient interface for these operations, however, these are not universal: some current e-infrastructures use emailed instructions or uploaded text files.

Some services may require distinctions between different members of a research group. For example, some members may be allowed to create new documents, some to amend existing ones, and some only to read them. For services that involve the consumption of assigned resources – such as equipment time, processor cycles or storage space – it may be desirable to assign quotas to individual group members. In these cases the group management platform and its interface to the e-infrastructure service need to express more than a simple binary member/non-member distinction: permissions and quotas may need to be associated with individual members.

Authentication Requirements

Whatever membership information is managed, and whatever interface is used, the ability to edit or view the group membership needs to be restricted, so the group management platform needs to be able to authenticate each PI and give them authorised access to their own groups. For most platforms federated authentication, whether using SAML assertions or X509 certificates, appears the best way to achieve this.

This in turn creates a requirement to link the PI's login to the e-infrastructure service with their login to the group management platform. Otherwise the PI, when logged in to the e-infrastructure service, will not be able to authorise that service to read the group(s) that they manage while logged in to the group management platform. Authentication systems with privacy as a high-priority design goal often include features designed to prevent different services identifying actions as being performed by the same individual. For example, the authentication system may release a different opaque identifier to each service, with only the identity provider being able to link the individual targeted identifiers. While it is possible to implement account linking services that let each user control which services they will allow to link together their activities, it is likely to be simpler to encourage identity providers to release the same ("untargeted") user identifier to the different components of an e-infrastructure. The REFEDs Research and Scholarship definition of services supporting academic collaboration[1] is designed to provide a framework for the release of such identifiers while keeping the privacy risk at an acceptable level.

Invitation

The same cross-linking issue arises for group members as for the group manager. If an e-infrastructure service is to use group information then it must be able to discover whether an authenticated user is indeed a member of the group entitled to use the resources they are requesting access to. The group information must therefore either contain, or be linkable to, the identifier by which each individual user is known to the service. While that could, in theory, be done using some other identifier shared between the service, the user and the PI, it appears most convenient (provided the linking issue discussed above can be solved) to have this done transparently by the group management platform. This requires that each group member must authenticate to the platform, at least when they first join a group, so that the platform can subsequently recognise users of e-infrastructure services as members of its own relevant groups.

In the unlikely event that a group consists only of existing platform members, the platform may already have the information required. In most cases, however, the PI will need to invite one or more of the desired group members to join the platform so they can be added to the group. Since this invitation would often be sent by e-mail, the platform may itself provide an Invitation function, accepting an e-mail address entered by the PI and sending that person an invitation to join the platform and the group. If the invitation e-mail contains a URL including a unique token then the recipient can accept the invitation by simply copying the URL into their browser (or clicking on it if their policy permits) and authenticating to the platform using their existing federated credentials. If joining the group requires a member to agree to additional conditions then this can be done through the same e-mail but in this case invitations need to be sent to all invited members, including those already known to the platform.

Using an e-mail address to identify the desired group member also solves the problem of getting the right "John Smith" into the group. E-mail address is almost certainly the truly unique identifier (unlike name!) that a PI is most likely to know for their collaborators. For this reason the Interfaces section below suggests that all group members be specified by e-mail address, even if the platform subsequently determines that there is no need to send a new invitation to those members who are already registered.

While most invitations appear likely to be initiated by the group manager wanting to share resources, in some cases a request for access (for example to a dataset) might be initiated by an individual researcher. This function could also be provided by a group management platform, though it would be up to the group manager to establish the requester's bona fides before accepting the request. As with the problem of identifying desired collaborators, this can only be done through real-world contact, not through the group management platform alone. The REMS system is an example where researchers apply and are required to commit to dataset terms and conditions.[2]

The flowchart on the following page illustrates how the joining process might appear to the group manager and the invited group member, as well as the logic followed by the group management platform.

[1]

1 refeds.org/category/research-and-scholarship/
2 terena.org/activities/tf-emc2/meetings/26/rems.pdf

Invitation to join a group

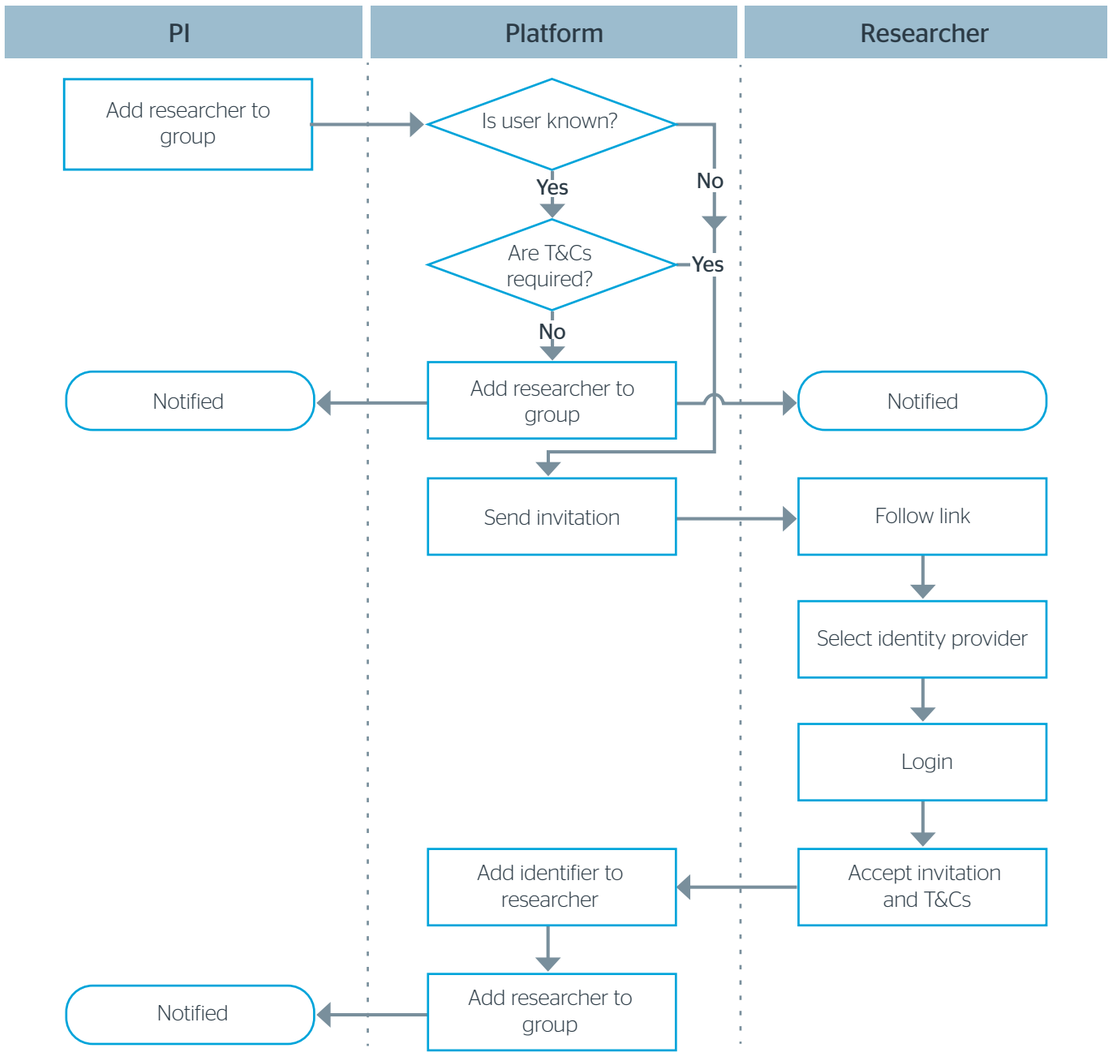

Alternative Authentication Sources

The discussion above assumes that each user will be able to provide an authenticated identity from a federation with which the group management platform and e-infrastructure services have a relationship (either being members of the same access management federation or through an inter-federation agreement such as eduGAIN[3]). However, as collaborative groups spread beyond the UK research community that assumption is less likely to be met. Researchers who are not associated with a university may not have a suitable login account; those from universities who have not joined a federation may not be able to make theirs available. Since it is undesirable that lack of an identity provider should prevent someone joining a research collaboration, the group management platform should consider what other sources of authentication might be available.

The obvious solution is for the group management platform to include a standalone identity provider that issues usernames and passwords for these unfederated users: sometimes referred to as a 'home for the homeless'.[4] Since these individuals have no organisation to take on the responsibilities of an Identity Provider (IdP) – vouching for their identity, helping them manage credentials and applying any necessary sanctions – the group manager/PI will need to fulfil those. However, the person who leads the research collaboration is probably the next best able to do that in any case.

An alternative approach is to use an identity provided by someone other than the individual's home organisation. This might, for example, be a social network service such as Facebook or Google, or a research community service like Umbrella.[5] These may offer some of the required identity management functions, in particular issuing, managing and verifying login credentials. However, as discussed in our paper on Federated Authentication[6], they are unlikely to be tailored to the needs of a particular group or service; functions such as enforcement of service or group policy are unlikely to be available.

Furthermore using a standalone or external IdP is unlikely to give the usability and security benefits of single sign on, so a fully federated identity is likely to have advantages for everyone.

If a group management platform does offer non-federated authentication then the e-infrastructure services that consume its groups are likely to need to accept the same source(s) of authentication. For a standalone IdP this means the platform must provide SAML- or certificate-based authentications to those services. If the platform accepts one or more external authentication sources then it may be technically simpler for the platform to act as a gateway between those sources and the services, rather than having each service separately implement the same sources of authentication (and possibly linking their targeted identifiers).

Whether federated, local, or external sources are used for authentication, the group manager/PI will need to ensure that these systems and processes satisfy the requirements of the relevant e-infrastructure services.

[1]

3 services.geant.net/edugain/Pages/Home.aspx
4 See, for example, DARIAH refeds.org/meetings/boftnc14/slides/DARIAH-AAI_FIM4R-Bof-TNC-v1.pdf
5 See, for example, EUDAT refeds.org/meetings/boftnc14/slides/EUDAT-jensen.pptx
6 bit.ly/1BmYJUj

Shared Management

As a group grows larger it may become impractical for a single individual to manage it, or to identify and vouch for all its members. In a research context a Principal Investigator may well wish to grant membership to "my colleague and her research students" or for a large international collaboration to have local staff manage participation from each country or organisation. This suggests that the group management platform should also allow a PI to grant selected group members the ability to add or remove individuals from membership.

More complex options for shared management, for example, elections to group membership, could be provided. However, complex schemes for sharing group management tasks will require complex interfaces and may create difficulties for accounting and accountability (see next section) that may not be acceptable to service providers. These facilities should be carefully designed to reflect user and service needs.

It appears, however, that any shared management function only needs to be implemented within the group management platform. Provided the platform can trace how membership, rights and resources came to be allocated to a particular user, this should not need to be communicated to the services that consume the resulting group memberships.

Accounting and Accountability

In several situations there may be a need to link activity on an e-infrastructure service to the user who initiated it. If a program run failed or required operator intervention then the operators may wish to contact its owner to explain the problem ("notification"); where use of the e-infrastructure is subject to quotas or charges the PI may need to track how those are being spent ("accounting") or to allocate limited (or costly) resources among users ("allocation"); if there is a breach of policy or other applicable rules then the responsible individual may need to be held accountable ("accountability"). Since the e-infrastructure service will often have only an opaque identifier for the user, linking that identifier to the user will often require support from the group management platform.

These processes should be designed in order to respect the privacy – both personal and professional – of individual users. In particular associating activities with named users should be a separate accounting/accountability process, not part of routine operations. Continually associating a name with every action is unnecessary and may increase the legal risks for services, platforms and identity providers, making them less willing to participate. For some datasets it may be appropriate to offer anonymous access, but with the possibility of invoking an accountability process if terms of use are breached. Thus the normal method of communication should be that a service asks the group manager to pass a message to the individual user, not that the group manager discloses the user's contact details to the service.

Since policies for access to e-infrastructure services and data may include particular accounting or accountability requirements, group management platforms should document the functions they offer. Although responsibility for policy compliance is likely to rest with the Principal Investigator, the availability of suitable accounting/accountability functions may be relevant when service or data owners decide which group management platforms to support.

Notification

Authentication systems are generally designed to be initiated by the user: they make a request to access to a service, prove their authorisation to do so, execute commands, and so on. If, instead, the service wishes to initiate communication with the user, this may not be supported. The simplest option is to ensure that a notice is displayed next time the user logs in, but this may not suit the access technology or the requirement (e.g. if the notice is information about scheduled unavailability of the service).

To provide more immediate communication, services may invite PIs or individual users to register an e-mail address to receive messages concerning their group or account. Allowing users to choose the appropriate address for notifications is preferable to requesting it from their identity provider because it allows messages be directed to a project or operations mailbox. As e-infrastructures develop to provide a greater number of services there may be an opportunity for group management platforms, which are likely to already know the user's e-mail address, to simplify this process by providing a mail forwarding service. This could, for example, allow a service to send notification e-mails to the platform, which then forwards them to the relevant user or group manager address.

Accounting

Accounting is the process of generating reports of what e-infrastructure resources were used. There may be many reasons for this, including management of projects or resources and billing. The key distinction from Allocation (see next section) is that Accounting does not directly regulate the consumption of resources. Since Accounting is generally not time-critical, it can be done in many different ways and with different degrees of automation. A group management platform could contribute to these in various ways.

The simplest form of accounting is done locally by e-infrastructure services keeping a record of the resources used by each account. Unless services have a direct relationship with their users such a record is, however, of limited use beyond local capacity planning. Where access has been granted to a group or project, reports of the use by members of that group are likely to be much more useful to the group manager. This requires the service to assign user sessions to a particular group, something that is best done using the group membership information obtained from the group management platform at the time each session was authorised. Although user activity could potentially be assigned to groups retrospectively, this is likely to be unreliable if group membership changes. Per-group accounting does not require any additional support from the platform beyond the existing authorisation exchange. This would, for example, allow Principal Investigators to be invoiced for the use their groups had made of the service.

Even if charging is based on a group's total use, a group manager may need more detail about individual use, for example if a user has consumed more than their intended share of the allocated resource. However, the e-infrastructure service can only report against its own account names, which may be the opaque strings provided by federated authentication. Accounting for individual use may require the service report to be merged with user data obtained from the group management platform. Where a group uses multiple services, these might be combined into a single report. Various tools exist for processing reports from federated systems, for example Raptor[7], developed alongside the UK Access Management Federation. A group management platform could offer various levels of support for reporting, from simply providing tools for manual use to automating the process of obtaining relevant logs (if necessary including cross-checks to ensure accurate billing) from services. Automation would, however, require services to adopt a common approach to providing log extracts via an authenticated download.

[1] 7 iam.cf.ac.uk/trac/RAPTOR/wiki/Software/Overview

As e-infrastructure use grows, it is increasingly likely that the same user will be a member of more than one group that uses the same service. This may raise complications for authorisation, accounting and allocation, since the e-infrastructure service may have no way to determine which group's resources to offer or, if charging or quotas are used, which group these should be assigned to. General purpose e-infrastructures, in particular, may need to offer users a choice, when they request access, between the various group memberships that could authorise their access. Although this choice will use group membership information from the platform, the interaction is likely to take place directly between the service and the user.

Allocation

Allocation is the process that ensures the resources available on an e-infrastructure are shared among its users, by placing either fixed or dynamic limits on the amount of resource that each group or individual user can consume. Unlike accounting, which reports after the event and is principally interpreted by humans, allocation is forward-looking and principally implemented by machines. Allocation decisions also need to be brought into force within a known (usually short) time window, so are more likely to be implemented through automated rather than manual processes. Again, there is a range of options, some of which may benefit from, or require, support from the group management platform.

Once an e-infrastructure service is able to obtain group membership information as part of its authorisation check, both per-user and per-group quotas can be implemented locally, without any additional interactions. This can include both static and dynamic enforcement, for example a 'fair-share' scheduler might lower or raise the priority of users or groups based on whether their recent activity has been above or below the target rate. As with accounting above, if a user is a member of more than one group on the same service, it may be necessary to let them select which group each activity should be assigned to.

Whereas quotas within a single service will generally be agreed in advance between the service and the group manager, more interaction will be required if quotas can be transferred between a manager's groups, or between different services. For example, if one activity requires less CPU time or money than expected, the manager may wish to re-allocate the spare resource to another subgroup or a different service. While this could be done by e-mails between the group and service managers, group management platforms might consider providing a resource management tool, though this would require standard protocols to be developed and supported by both the platform and the relevant service(s).

Accountability

Where a user has broken the terms of an acceptable use policy or other conditions applying to their use of the research data or service, the individual may need to be held accountable. Since the service where the breach took place may only have an opaque identifier for the account, the group management platform will need to be involved to associate the misuse with the responsible individual. This requires both the technical ability to attribute activities to individuals, and also policy agreements between the service and platform covering how problems will be reported and what actions will be taken.

For small groups, particularly where direct e-mail invitations were used, it is likely that the Principal Investigator will know the identities of all group members and be able to contact them directly. Where management of larger groups is shared with others it may be more appropriate to have them also deal with situations where accountability is required. As discussed in the Federated Authentication paper[8], a person or organisation local to the user is more likely to be able to impose appropriate accountability measures or sanctions than a remote service or group manager. Accountability policies and processes should therefore be designed to take the problem report to the responsible individual, rather than disclosing the individual's identity to a service manager who is unlikely to be able to use it to impose sanctions in any case. The UK Access Management Federation's Rules of Membership[9] already require those home organisations that declare the ability to identify users to also deal with any breaches of service policies by those users.

While services may be able to take technical action to block access by an individual or group if a problem is not dealt with, this is best kept as a reserve power for extreme situations where the service can no longer trust the group manager and home organisations to meet their agreed responsibilities to protect services and data. Such action must be communicated to both the group manager and home organisations, as if they are not made aware of the reason for a technical suspension they may well "solve" the user's access problem by simply changing the allocated identifier, leaving the service and data once again at risk.

[1]

8 bit.ly/1GACEYa
9 bit.ly/1cXL7bZ

Group Management within an E-infrastructure

The diagram shows how these various functions of the group management platform interact with the other components of the e-infrastructure.

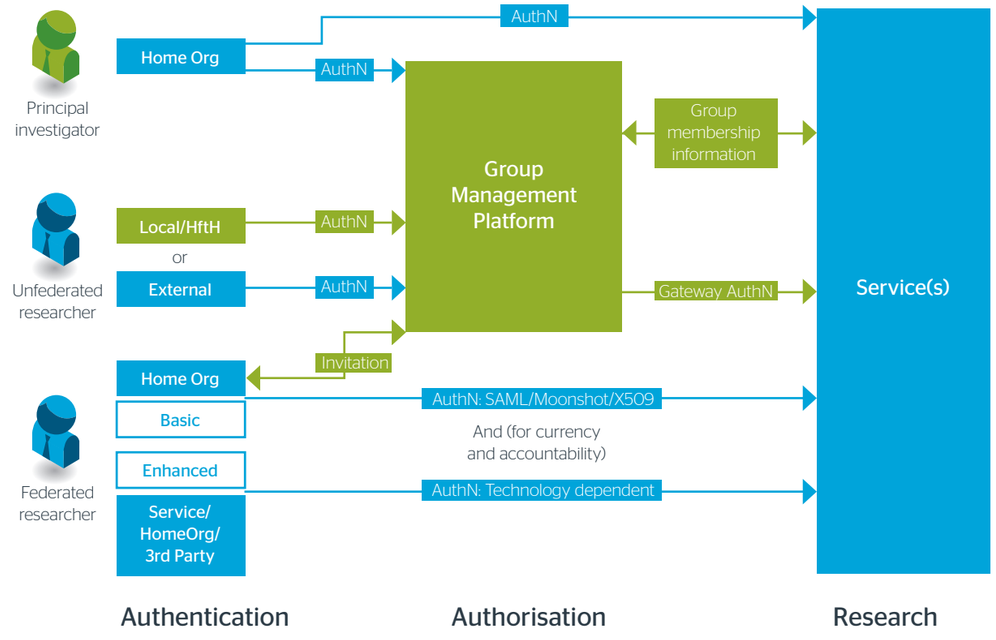

Location

Most current group management platforms are provided as part of a specific e-infrastructure service. This ensures that the group management functions satisfy the particular requirements of that e-infrastructure.

Having the two functions subject to the same management should ensure they implement consistent policies and may help with some of the challenges of linking identifiers. Where an e-infrastructure consists of multiple services – as for example with DiRAC or EGI – the same group membership information may be distributed to the different services within the group.

However, treating group management as part of a single e-infrastructure is likely to create significant difficulties for research that needs to span multiple infrastructures, or that involves different research communities, or simply where a PI wants to work with the same group of collaborators on a different e-infrastructure service (for example to use a video-conferencing or collaboration platform for meetings between a group of e-infrastructure users). At present this may well require re-registering each user on every e-infrastructure and reconstructing the groups on each one's management platform. Systems for incorporating researchers from outside existing research and education federations also need to be re-implemented separately by each e-infrastructure that wants to provide them.

With requirements for cross-discipline and cross-infrastructure research likely to increase in future there will be a growing need to share group management platforms, and the information and services they provide, between e-infrastructure services. This is likely to involve existing platforms making their services and information available to e-infrastructures other than those with which they are directly associated. In future there may also be benefit in offering standalone group management platforms to increase the range of facilities that can be provided across the UK and international research infrastructures, and to make more efficient use of the effort that is expended on creating and supporting them.

If group management platforms are to communicate with multiple e-infrastructure services, this will require greater adoption of common standards for that communication by both services and platforms, since configuring each individual link is unlikely to be feasible. As is discussed in the next section, at present this communication is largely contained within e-infrastructures and a remarkably wide range of different approaches has been adopted. Although it would be desirable in the long term to move towards standard protocols for all these communications, significant progress could be made by adding support for those protocols to existing systems. This would allow e-infrastructure services to obtain group membership information from group management platforms other than their own, as well as allowing groups already existing on those platforms to be used on other services or for cross-infrastructure research.

If services are going to rely on platforms other than their own for group management information, some work will also be needed to ensure that policy requirements are satisfied. Existing federation policies and inter-federation agreements may well provide a basis for these and any additional agreements should be standardised as far as possible. This should simplify policy discussions for identity providers, service providers and users by reducing the number of agreements each of them needs as members of multiple communities, federations or infrastructures.

Interfaces

As a central component of an e-infrastructure, the group management platform must interface with a number of other systems. In particular there is a human interface to the Principal Investigator and technical interfaces to identity provider systems and e-infrastructure services.

This section looks at the requirements for each of these interfaces, what is used at present, and the challenges of increasing inter-operability if e-infrastructures are to become less tied to dedicated group management platforms.

Principal Investigator Interface

The main user of the group management platform is likely to be the Principal Investigator who uses it to create and manage group(s) of collaborators. Since this is a human/machine interface it does not require formal technical specifications in order to work, though some general principles can be identified to make platforms easier to use.

A wide range of interfaces appear to be offered to group managers by present systems, from graphical interfaces presented through a web browser to simple e-mail messages sent to system administrators. Relying on the source address of an e-mail to authenticate the PI who is entitled to change group management does not appear a particularly secure or scalable solution; cryptographic authentication – whether by signed e-mail, X509 certificates or federated authentication – is likely to be a minimum requirement in future.

The interface needs to let the group manager review the current members of the group and the permissions they have been granted. However, providing a list of all users known to the platform, in case the manager wishes to add some of them to the group, is unlikely to be satisfactory from either usability or privacy perspectives. If group management is successful then the list of registered users should soon become too long for convenient use; there is also a growing risk of confusion among researchers of the same name, particularly when the individual the PI wants to invite is not registered with the platform but someone else of the same name is. Since identifying the correct individual can only be done outside the platform in any case, it is better for group managers to enter a unique identifier that they know refers to the right individual (for example an e-mail address exchanged in person) and for the platform to either add an existing user to the group or to send a new member invitation, registering the user with the group and the management platform as described above.

Although there is no formal requirement for group management platforms to offer the same interface, it will clearly be simpler for Principal Investigators if they do not need to learn a completely new interface for every new e-infrastructure they use. E-infrastructure services that can accept group membership information from different platforms will offer a significantly improved experience for PIs.

Identity Provider Interfaces

The group management platform will need to authenticate all its users: group managers so it can authorise their access to the correct groups, and group members so it can link their group memberships to the identities they will use to authenticate to e-infrastructure services. Although a group management platform could in principle provide authentication services for all of its users, as discussed above and in the Federated Authentication paper it appears better to use federated authentication services wherever these are available.

Within research and education federations three main protocols are used for authentication: X.509 certificates, SAML/Shibboleth, and RADIUS/Moonshot. At present there are e-infrastructures in the UK that use at least the first two of these for group management and the third is being piloted. In principle there appears to be no reason why the group management platform and the e-infrastructure service must use the same authentication protocol, so long as they provide either the same identifier for the user or else identifiers that can be linked. Thus, for example, a group management platform that offered a web interface might use SAML/Shibboleth for authentication and provide the resulting group information to an e-infrastructure that used Moonshot to authenticate SSH access by authorised users. Or if a certificate-based e-infrastructure used SAML authentication to unlock a certificate proxy then the group management platform might use the same SAML protocol for its own authentication, while the e-infrastructure services continued to use the X.509 certificates containing the same identifiers.

The group management platform's separation of the authentication and authorisation processes provides a significant simplification of the identity provider interface, since all the platform needs to receive is a unique, persistent identifier for each authenticated user. All the other information about the user – their group memberships, name, e-mail address, etc. – are either provided by the group manager as part of the process of joining the group or may already be known to the platform as a result of membership of other groups. Thus the platform only needs to be concerned with the protocol required to receive that identifier value: there is no need for protocols or agreements to assign any meaning to the identifier. If the identity provider cannot ensure that the same value will be provided to the group management platform and the e-infrastructures it serves then the platform may need to provide an account linking service for each user to associate the different identifiers. A number of different protocols for account linking have been developed, though, as far as is known, none are currently used by UK e-infrastructures.

If, as suggested above, the group management platform also wishes to accept authentications from services outside federations – for example, social or community authentication providers – then this may require support for a wider range of authentication protocols. Such a platform might also provide a gateway function for e-infrastructure services so these do not all need to implement these external protocols.

Service Provider Interfaces

The group management platform also needs to communicate with e-infrastructure services, to provide them with the group membership information that they use to decide what actions each authenticated user is authorised to perform. This interface requires agreed protocols at two levels – the technical means by which messages are passed, and the semantic meaning of those messages when they are sent and received.

Since many current group management platforms are dedicated parts of specific e-infrastructures, there is a very wide range of methods used to communicate group membership information.[10] This variety even extends to the time at which membership information is passed: in some cases membership changes are notified to the service at the time they are made, or updated on a regular (e.g. daily) schedule; others allow a user to request a signed long-lived token asserting that they were, at the time of the request, a member of a particular group; in others the user's group membership may be checked either at the moment they log in to the service or when they request an action where the permission to perform it depends on membership of a particular group. Systems where membership information is transmitted proactively to the service are referred to as "push" (with certificate-based systems the information may be pushed to an attribute certificate that the user later presents to the service). These run the risk that information may be out of date by the time the user requests access to the service. Conversely "pull" systems allow the service to request up-to-date information at the instant it is required for an authorisation decision, but make the availability of the group management platform as critical to the e-infrastructure as the authentication systems are. "Pull" systems can also limit the information they request to only that required for the current authorisation decision; "push" systems may need to provide all the information they have about the user in case it may be relevant to authorising a future action. RFC2904[11] describes these different authorisation architectures. All systems need to make a policy choice whether or not to authorise access if current information is not available.

The format in which information is passed also varies widely – from a complete list of all memberships (as represented in a unix /etc/groups file) to mere confirmation that user X is indeed a current member of group Y. This variety is likely to be the biggest barrier to connecting different e-infrastructures to the same group management platform. Where standalone group management systems have been implemented, their operators have often had to configure each new service individually before making it available through the group management platform. However, there are signs of progress, with a recent survey by the GN3+ project[12] finding that most platforms offer support for the SAML2 Attribute Authority query protocol.[13] Service providers that implement this protocol, at least in addition to their native method, are most likely to be able to obtain information from external group management platforms.

As well as agreement on the technical protocols used to transfer group membership information, moving towards a common service provider interface requires agreement on the meaning given to the content of those messages. The simplest approach, used by many current e-infrastructures, exchanges only an opaque identifier for a user and an opaque identifier for the group. This confirms that the user is a member of the group. However, the service provider alone decides what membership of that group means: what permissions and quotas are associated with it. Principal Investigators must discuss their authorisation requirements with each service they use, identify the group names that will be used to apply authorisations to users, and then create and allocate those groups manually on the group management platform. To use the same group on a different service, either the service needs to set up a group of the same name with the relevant permissions, or else the group management platform needs to be told to use different names for the same group depending on which service is making the request. In either case the platform, which is supposed to be the PI's main interface, is limited to adding or removing users from groups since it does not know what individual group memberships signify.

[1]

10 See, for example, refeds.org/meetings/may14/slides/aa-refeds-201405.pdf
11 tools.ietf.org/html/rfc2904
12 docs.google.com/spreadsheets/d/1SZGw3NI9OOTToiuHibv3b0zKteU-xlaMj1QNm225J78/edit#gid=0
13 bit.ly/1RkHXyf

A better division of functions between service and platform might be possible if roles within groups are defined, for example the Jisc Community site[14] allows each user to be assigned either "administrator", "editor" or "contributor" status within a group. This would allow the group management platform to represent these roles as having different privileges within the group, and might also make it easier to share a group (which might be named after a project) between multiple services, with each service applying its natural meanings of privileges to each role. For example, on a storage infrastructure the "contributor" role might be able to read and amend existing files, the "editor" role to additionally create new files, and the "administrator" role to create new directories; on a video-conferencing infrastructure a "contributor" might be permitted to participate in a call, but only an "editor" could create a new conference. Using a group on another service should then simply be a matter of informing the service of the group name and the platform where membership information can be obtained, and authorising the service to access group management information from the platform. Individuals would then inherit the appropriate rights for their role within the group, without those having to be re-defined and re-assigned on every platform.

Some research groups may need more fine-grained access control, perhaps even managing the rights of individual users to access individual files from particular locations. Others may wish to manage resources through quotas, for example, of disk space or CPU time, rather than rights over individual files or directories. This is likely to require the exchange of significantly more complex messages than just group membership, with only a limited range of services being able to implement the particular authorisation rules required. It is not clear whether this can be done in a sufficiently general way to be provided through a general-purpose group management platform.

Whatever level of meaning is assigned to group membership, some system is required to ensure that group names are unique. Otherwise if two projects each create a group called "administrators" a service could easily confuse them and give the project members unintended access to each other's resources. Group management platforms can, and should, ensure that their names are locally unique, but where a service can accept groups from multiple platforms this not sufficient, as two platforms may inadvertently create the same local name. Fully-Qualified Domain Names (FQDNs) are normally regarded as uniquely identifying Internet services, so services should normally incorporate the FQDN of the platform into the label they use as the group identifier to avoid clashes. Where a project does intend group names from different platforms to be treated as equivalent, this must be agreed between the relevant services, platforms and group managers.

Inter-operability

At present nearly all e-infrastructure services are closely coupled to a single group management platform, indeed the platform may form part of the service. Each service and platform has tended to support only a single authentication technology – either certificate or SAML federation – though some e-infrastructures provide additional options and the use of SAML authentication to certificate stores can permit some cross-technology operations.

This has disadvantages for all those involved: users may need to obtain multiple credentials to access different e-infrastructures; PIs need to re-define their groups if they wish to use more than one e-infrastructure (and cross-infrastructure projects may be impossible); home organisations need to make separate technical and contractual arrangements for each e-infrastructure their users wish to use (many will not have sufficient resources or skills, so smaller user groups risk being excluded for lack of effort); every e-infrastructure needs to start essentially from scratch if it wishes to support new user communities or authentication protocols as each e-infrastructure's effort only helps itself.

For e-infrastructures to support cross-community and cross-infrastructure working it appears essential that group information be obtained from multiple sources; treating group management as a separate "PI interface" function should also simplify the current challenges for home organisations, e-infrastructure services and users. The key development towards achieving this inter-operability appears to be for management platforms and e-infrastructure services to support some common method for transferring basic group membership information. Initially, at least, this could be provided in addition to whatever means is currently used to add group privileges to the e-infrastructure. This would at least allow groups and individual users to be imported from elsewhere.

Alongside work on technical exchange of group membership information there needs to be consideration of whether any new policy agreements are required, as e-infrastructure services will be relying on external group management platforms, whether run by other e-infrastructures or independently, for some of the assurance that has previously been provided by their own processes. Group management platforms can be viewed as acting as service providers (recipients of authentication information) towards home organisations, and identity providers (providers of authorisation information) towards service providers, so most or all of what is required should be available under existing federation agreements. Where service providers require specific assurances of process or responsibility, agreements may be needed on how these, or suitable alternatives, can best be provided.

Increased inter-operability may in time lead to a focus on a smaller number of group management platforms or software implementations, with benefits for all their users. Development efforts on the security, functionality and performance of these systems will benefit a wider range of e-infrastructures; standalone group management platforms could be deployed in highly scalable ways through virtualisation and DevOps approaches.

[1] 14 community.jisc.ac.uk/

Conclusions

The implementation of features by current group management platforms appears to follow one of two distinct patterns. Some features, such as invitation, are already implemented by several different platforms using largely similar approaches.

These are generally features internal to the platform, so whether or not a particular platform implements it depends on the particular requirements of its user community. This suggests that when other platforms do require these features, it should be relatively straightforward to add them following those existing models. Working code may even be available.

However, for many external interfaces, such as that to the service provider, there is much less consensus on the right approach. Here development may be fragmented across a number of different platforms or even suppressed entirely through lack of critical mass. Where this causes duplicated work or incompatibilities between systems it appears that investigations or pilot studies, by going beyond the immediate needs of individual systems or infrastructures, might help to clarify how an e-infrastructure of interoperable components might be achieved. This could establish targets for currently incompatible infrastructures to develop towards.

This section first identifies the internal features where clear models exist that individual infrastructures can adopt as and when their requirements develop. It then considers in turn each of the more complex interfaces where additional action may be required to help move towards interoperability.

As the range of possible collaborations increases, it will become more important for services and group management platforms to document the policies and processes they support. This will allow researchers, services and data owners to identify suitable e-infrastructure components to support their work. Where possible these policies and processes should adopt existing federation or community standards to facilitate inter-working.

Internal Features

Several features of group management platforms are internal to each platform and do not have significant implications for other components of the e-infrastructure. Individual platforms can implement systems for invitation, local authentication and shared group management as and when their users require them without affecting their technical or process compatibility with other e-infrastructure components. Such systems may, however, affect platforms' ability to comply with others' policies, for example on accountability or assurance of identity. Such policy requirements should be borne in mind and the conclusions documented as part of the platform's description of its functions. Examples of each of these functions already exist in national and international e-infrastructures using approaches that could be followed by those wishing to implement them.

Invitation/Joining

It appears that most e-infrastructures use one of two processes for joining a group: either a known person is invited to join by the group manager, or else a previously unknown person requests membership from them. In each case some manual process will be needed, outside the group management platform, to ensure the intended individuals are granted membership of the group. However, group management platforms could provide more technical support for the joining process. Offering a group manager a complete list of individuals known to the platform is unlikely to be the best option: interfaces do not scale well as the number of members grows, there may be privacy issues if a platform supports groups in different areas, and there is a risk of confusion where members have common names. An invitation approach such as that provided by SWITCH[15] where the manager enters a known identifier for the intended group member (typically an e-mail address) avoids these problems. Where individuals need to apply for membership, a workflow process is likely to be appropriate. This should capture sufficient information about the applicant for the manager to be able to check whether they are entitled to join the group; it may also, as with REMS[16], be used to have the applicant agree to terms and conditions of use.

Local Authentication/Home-for-the-Homeless

Where a group management platform needs to include a small number of users from outside the academic research community, a 'home-for-the-homeless' may be an effective way to provide them with a username and password to access their group resources. Such a service may also provide attributes that are not available from home organisations, but these are likely to be self-asserted by users so unsuitable for making authorisation decisions. Single-purpose login accounts are, however, inconvenient for both the user and the group manager: a federated identity that can provide single sign on and benefit from the home organisation's account management and policy enforcement is preferable. For large numbers of unaffiliated users, interfaces to third party authentication sources (described below) may be more suitable. Group management platforms offering homes-for-the-homeless need to ensure that their security policies and practices match those expected of federated identity providers or else these accounts will reduce overall trust in the platform. The DARIAH project has an example of a local home-for-the-homeless.[17]

Distributed Group Management

Distributed group management is likely to be required by most platforms supporting groups of any significant size. A wide range of implementations already exist (for example those of SURFconnext[18] and VOMS[19]), with others proposed.[20] The choice of which to offer is likely to depend on the requirements of the services and users that each platform wishes to support. Group management practices can also affect the security and policy assurances that platforms can offer to services, so these should be documented.

Accountability

Group management platforms will require some accountability policy and process, if only to suspend access to the platform by those who misuse it. Most access management federations will also require members (including platforms) to assist others in dealing with incidents: as an intermediary between services and home organisations a platform should at least be able to pass on a complaint from a service provider to the appropriate group manager and home organisation. Additional requirements for notification, accounting, allocation and accountability appear to vary considerably between existing communities, though there is a general trend to require more in these areas. Group management platforms should regularly review the requirements and plans of the communities they wish to serve.

[1]

15 bit.ly/1J2OPNU Page 9
16 terena.org/activities/tf-emc2/meetings/26/rems.pdf
17 bit.ly/1RkHXyf
18 bit.ly/1LDvxO7
19 wiki.egi.eu/wiki/FAQ_VO_Management
20 bit.ly/1cXWktg

Interfaces

Third-party Authentication

As use of e-infrastructures expands beyond the academic research community, an increasing number of potential users will not have an existing relationship with either an identity federation or a registration authority for recognised digital certificates. Although a home-for-the-homeless function may be able to issue these users with a username and password, this is a relatively onerous process for the group management platform and inconvenient for the user. Platforms that wish to support significant numbers of un-federated users are likely to find it more efficient to use third party sources of identity, either within the research community (e.g. Umbrella[21]) or beyond (e.g. social network services). Third party authentication sources may also offer features such as multi-factor authentication that are not yet widely available through federations.

Although a few e-infrastructures have investigated and used un-federated authentication sources, this seems a sufficiently complex but widely-applicable area that more general recommendations would be helpful. These could, for example, cover which sources of authentication are most likely to be useful to e-infrastructures and their users, how to incorporate these into group management platforms (for example should the platform provide an authentication gateway for connected services), and any policy issues that need to be considered. Documented pilot studies would help other e-infrastructures implement third-party authentication as their user communities expand.

Cross-technology Authentication/Authorisation

Current e-infrastructures divide into two groups: those that use X.509 certificates for authentication and those that use federated login. Although it is possible for a federated login user to translate this to a certificate using proxy services such as MyProxy[22] and SARoNGS[23] it is not clear how a group containing users from both technologies might be managed or implemented or how the different policies might be reconciled.

A further challenge occurs if different authentication technologies are used by the group management platform and the service for which it manages authorisation: for example if the platform has a web interface using SAML authentication but the service uses non-web Moonshot. In theory, since both originate from the same authentication source, both protocols should be able to carry the same unique identifier for the user, but the technical and configuration issues have not yet been investigated.

In each case documented pilot studies would be very helpful in determining the best approach and assisting others to adopt it.

Service Provider Interface

The key component in expanding access to e-infrastructures will be the interface between the group management platform and the service providers to which it provides authorisation information. For groups to have access to multiple services, and services to be able to accept group membership information from multiple platforms, different platforms and services must support the same method for communicating membership messages.

At the protocol level SAML Attribute Assertions appear increasingly recognised as the common standard that should be supported. E-infrastructure services and group management platforms should be encouraged to support this protocol alongside their existing methods for communicating group membership.

Using multiple sources for group membership creates the risk that two different sources will use the same name for a group, possibly providing unintended access to information or services. Services that accept group membership information from external sources should normally treat the fully-qualified domain name of the source as part of the group name, thus ensuring that, for example, "example.ac.uk:secret" and "example.edu:secret" are treated as different groups. Where a single group is obtained from two different sources, an appropriate naming convention should be agreed between the group sources and the service provider(s).

Group membership messages are typically simple, perhaps containing only an identifier for a user and an identifier for a group. Some systems offer additional information about the type of membership (e.g. "contributor", "editor", "administrator") rather than setting up different groups for each of these roles. If such roles are common to a number of different e-infrastructure services then support for them might reduce the amount of negotiation required when connecting an existing group to a new e-infrastructure. It has also been suggested that messages might carry more detailed information about assigned quotas or other rights but it is not clear whether these can be generalised or if they would always be specific to individual services. Further investigation is needed to determine whether common role definitions would be useful across e-infrastructure services.

Summary of Recommendations

Where e-infrastructure services and group management systems wish to implement new internal features – such as Invitation. Local Authentication, Distributed Group Management and Accountability – they should consider adopting existing approaches, both to benefit from work already done and to make their systems more familiar to users and group managers who have experienced the same functions elsewhere.

E-infrastructure services should consider implementing standard interfaces to facilitate access to group information from other services and platforms. Services investigating, or providing, interfaces to third party or cross-technology authentication should be encouraged to document these studies to allow common approaches to be adopted.

Services and group management platforms should document the policies and processes they support, and consider building on existing federation or community standards, in order to facilitate inter-operability between e-infrastructure components.

[1]

21 bit.ly/1EwTm4H
22 grid.ncsa.illinois.edu/myproxy/
23 cts.ngs.ac.uk/

Share our vision to make the UK the most digitally advanced education and research nation in the world

jisc.ac.uk

Jisc
One Castlepark
Tower Hill
Bristol, BS2 0JA
0203 697 5800
info@jisc.ac.uk