A hand-drawn diagram of a rocket launch on a grid background. The rocket is a simple line drawing with a central body and two side boosters. A dashed line indicates the vertical path of the rocket. A circular target symbol is drawn in the upper right quadrant. In the lower left, there are some mathematical expressions: eV/I and $I \cdot \ln\left(\frac{M_1}{M_2}\right)$. The background is a light blue grid with white dashed lines and small 'x' marks at the intersections.

A Proposal for a Common Global Access Management System for Education & Research

Josh Howlett
JANET(UK)

9 May 2011

1. Executive summary

In recent years, user requirements have driven the development of a set of access management systems to support inter-organisational activities within the global research and education community. These systems have, in general, addressed separate and often quite different application domains. Viewed independently, these access management systems have generally been highly successful. However these systems have, in general, addressed separate and often quite different application domains and, when considered as an ensemble, two issues become apparent:

- Their multiplicity imposes significant complexity and costs on organisations and users, by requiring them to interact with a number of dissimilar access management systems.
- Despite their multiplicity, these systems fail to address the inter-organisational access management requirements of many applications, resulting in significant opportunity costs.

An access management system that not only provided an access management solution for these other use-cases, but also for those of existing applications, would clearly be highly desirable. A new access management technology, ABFAB, developed primarily by the education and research community and undergoing standardisation within IETF, may be able to deliver this.

The technology has already been substantially implemented by Project Moonshot, a JANET(UK)-led initiative in partnership with the GEANT project and others; the planned work will be completed during Q3 2011. The technology has been demonstrated with a number of applications; *little or no software modifications were required*.

This paper argues that an access management system that delivered value to significant parts of the global research and education community could be implemented within months, and at relatively little effort and cost, and without 're-inventing the wheel'. This would be achieved through re-use of the global RADIUS infrastructure that currently supports eduroam: it presently incorporates over forty countries on every continent, connecting thousands of organisations and many millions of users. Some minor modifications to this infrastructure would be required but these could be managed almost entirely through upgrades to existing software. Most of the required effort would need to be directed at modifying existing the RADIUS infrastructure's policies to address the new applications.

This access management system could confer a number of benefits to the community, including:

- Lower operational costs for service providers and their customers.
- Increased opportunities for collaboration and revenue generation.
- Improved user experience leading to greater adoption and use of services.

The purpose of this document is twofold:

1. To brief the reader of a proposed strategy for *establishing a common global access management system by April 2012*, that will hopefully lead to discussion and consensus. This paper does not seek a single access management system; it is instead advocating a common access management system that is available to those that wish to use it.
2. To ensure that the community is aware of the potential opportunities and other consequences of the technology, so that informed action can be taken.

There is no Internet presence for this proposal at present; the manner in which this proposal is taken forwards – if at all – is a matter for the community. However, a temporary mailing list has been created to support co-ordination in the interim. Parties who are interested in tracking and discussing this proposal are encouraged to subscribe to cgams-announce@jiscmail.ac.uk at <http://www.jiscmail.ac.uk>.

Acknowledgements

This proposal builds on the work and inspiration provided by many within the education and research community and elsewhere, and the support provided by the author's colleagues at JANET(UK). The author expresses particular gratitude to the following: Scott Cantor (Ohio State University), Dr. Antonio F. Gómez-Skarmeta (University of Murcia), Sam Hartman (Painless Security LLC), Luke Howard (PADL Consulting Pty Ltd), Dr Jens Jensen (STFC), Leif Johansson (NORDUnet), Dr Daniel Kouril (University of Masaryk), Dr Diego Lopez (RedIRIS), Dr Gabriel López Millán (University of Murcia), Linus Nordberg (NORDUnet), Alejandro Perez Mendez (University of Murcia), Dr Rhys Smith (University of Cardiff), Dr Ian Stewart (University of Bristol) and Klaas Wierenga (Cisco Systems Inc). Needless to say, the opinions expressed in this paper are entirely those of the author.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 238875, relating to the project 'Multi-Gigabit European Research and Education Network and Associated Services (GN3)'.

eduroam is a trademark of TERENA.

Table of Contents

1	Executive summary.....	1
2	Introduction	5
3	Vision.....	7
4	Considerations.....	9
5	Use-cases	11
6	Conceptual model	17
7	Proposed implementation	19
8	Recommendations & roadmap	23
9	Appendix I – ABFAB overview	25
10	Appendix II – Technical feasibility.....	29
11	Appendix III – Expected modifications to RADIUS infrastructure ...	33

2. Introduction

The emergence of access management as a strategic requirement for the education and research community is a relatively new phenomenon. Traditionally, ICT services have been organised, delivered and consumed within individual organisations. Where services were delivered externally, or consumed from external sources, it was generally possible to engineer *ad hoc* approaches that, while technically unsophisticated, satisfied the business requirement.

Over time various developing requirements began to exceed the capabilities of these approaches. These included:

- Secure authentication of correspondents and data confidentiality. These properties are necessary to support a broad range of activities, from serving web pages to delivering email.
- Management of access to, and ease of use of, the data network and its services and content to support “anywhere, anytime” education and research.
- The emergence of large-scale computational services, such as Grids, High Performance Computing clusters and Clouds, which are characterised by distributed users and/or infrastructure.
- The evolving regulatory and statutory environment, such as the obligations imposed by legislation relating to data protection, freedom of information and so forth.

The response by the organisations that support the education and research community’s ICT requirements has, in general, been successful. Notable achievements include the International Grid Trust Federation, eduroam, the identity federations and the TERENA Certificate Service.

These successes, however, have yielded two substantial challenges:

1. Viewed independently, these services provide a consistent experience to their customers. However, these services use a variety of different security technologies. Customers of these services, and the organisations that deliver them, must necessarily maintain expertise and infrastructure for each technology. This technical redundancy imposes costs on those organisations that must deliver and consume these technologies, and complexity for their users who must learn to use them.

2. Despite the diversity of available technical approaches, there are still many applications and services that cannot be used effectively between organisations because no appropriate technical mechanism exists. This lack of inter-organisational functionality imposes a tremendous opportunity cost by preventing these applications and services from enjoying the benefits of inter-organisational operation.

There appears to be growing consensus that these issues are becoming increasingly acute. The purpose of this paper is to outline a *common* and *global* access management system that may comprehensively address these:

- *A common system*, meaning that is an appropriate solution for a broad range of application and service technologies.
- *A global system*, meaning that it is accessible and usable for diverse, geographically or otherwise, user and service provider communities.

This proposal is based on the ABFAB architecture and technology that is presently being standardised within the IETF and implemented by Project Moonshot, a JANET(UK)-led initiative in partnership with the GEANT project and others.

This paper focuses on articulating a vision for a common global access management system, and a proposed implementation strategy; it is not technical documentation. Technical architects and specialists that require more detailed information to evaluate this proposal are directed to the appendices. Appendix I provides an overview of the ABFAB technology; Appendix II provides a summary of the main results of Project Moonshot, which is developing an ABFAB implementation; and Appendix III discusses the deployment implications of this proposal.

3. Vision

This paper seeks the establishment of a common global access management system that is capable of supporting every service and application used by the education and research community.

Figure 1 below provides a high-level depiction of this vision.




Figure 1

The bottom layer represents the user groups that comprise the global education and research community. The top layer represents the services and applications that need to be accessed by these users. These service and applications may be operated by the same organisation that a particular user is affiliated with, or a different organisation; there is no distinction between intra- and inter-organisational operation.




Figure 2

The middle layer, illustrated in Figure 2 above, represents the access management infrastructure that mediates the users' access to these services and applications. This infrastructure consists of brokers, which mediate *vertically* between users and services, and an interconnection fabric which mediates *horizontally* between brokers.

4. Considerations

This section discusses the most pertinent and practical implications of this proposal.

4.1 Scope

The system advocated by this paper could, in principle, support many or even most of the access management requirements within the education and research community. However this may not necessarily be practical or even desirable in many instances:

- there may be no business case (the costs may be too high, the user demand too low, etc)
- there may be pre-existing commitments and plans that cannot be changed
- there may be value in maintaining a separation between access management systems.

Thus, this paper is not advocating a *single* access management system; it is instead advocating a *common* access management system that is available to those that wish to use it.

4.2 Applicability

Those communities that wish to take advantage of a common global access management system will, in some cases, need to consider transition or co-existence with other systems. This may, for example, entail modifications to software or operational practices and policies.

Therefore, while this paper claims that such a system could be implemented in months, it does not argue that it will be useful to everyone within the same timeframe – if at all, for the reasons described previously.

4.3 Preparation

For many user communities, however, it won't be necessary to wait for someone to build a global system. As described in sections 7 and 10, the software and infrastructure necessary to realise certain use-cases already exists.

Consequently even in the absence of co-ordinated action, it is quite likely – if not inevitable – that this technology will get deployed and used. If left unchecked,

however, this kind of *ad hoc* deployment could result, at best, in a less effective system than one which had been managed in a more co-ordinated manner from the start; and at worst a system with the potential for significant security issues (section 11 provides some relevant discussion).

This paper strongly recommends that those stakeholders with a responsibility for delivering or consuming access management services develop a strategy for managing the deployment of the technology within their constituencies. A number of user communities have already expressed interest in using the technology and they may not be indefinitely patient. Therefore, this paper urges these stakeholders to act swiftly to ensure that the technology is appropriately deployed, to the benefit of all.

5. Use-cases

(This section draws substantially upon the IETF document “Application Bridging for Federated Access Beyond web (ABFAB) Use Cases” located at: <http://tools.ietf.org/html/draft-ietf-abfab-usecases-00>.)

This section briefly describes a selection of the user communities and use-cases that could benefit from the existence of a common global access management system. These are:

1. Cloud services
2. Grid infrastructure
3. High performance computing
4. Web content and applications
5. Roaming network access
6. Multi-domain network management and configuration.

These use-cases have been selected for discussion in order to demonstrate the breadth and diversity of use-cases within the community and their requirements. Many other use-cases have also been discussed in other venues but have been excluded from this discussion in the interests of brevity.

5.1 Cloud services

Many organisations are seeking to deliver services to their users using ‘Cloud’-based providers. This is typically motivated by a desire to avoid management and operation of commodity services which, through economies of scale and so forth, can often be delivered more efficiently by such providers.

Many providers already provide web-based access using conventional authentication mechanisms; for example, to ‘web-mail’ applications. The use of federated authentication enables organisations that consume cloud services to more efficiently orchestrate the delivery of these services to their users.

Frequently, however, users will prefer to use desktop applications that do not use Web protocols. For example, a desktop email client may use a variety of non-Web protocols including SMTP, IMAP and POP. Some cloud providers support access to their services using non-Web protocols. However, the authentication mechanisms used by these protocols will typically require that the provider has access to the

user's credentials. Consequently, the provider will require that users' credentials are regularly synchronised from the user organisation to the provider, which may have obvious implications for security and privacy, or else be provisioned directly by the provider.

The latter approach may be acceptable in the case where an organisation has relationships with only a small number of providers, but may become untenable if an organisation obtains services from many providers. Consequently organisation with a requirement to use non-Web protocols would prefer to make use of the credentials that they have already provisioned their users with, and to utilise federated authentication with non-Web protocols to obtain access to Cloud providers.

5.2 Grid infrastructure

Grids are large-scale distributed infrastructures, consisting of many loosely coupled, independently managed, and geographically distributed resources managed by organizationally independent providers. Users of Grids utilize these resources using Grid that allows them to submit and control computing jobs, manipulate datasets, communicate with other users, etc. These users are organized into Virtual Organizations (VOs); each VO represents a group of people working collaboratively on a common project. VOs facilitate both the management of their users and the negotiation of agreements between their users and resource providers.

Authentication and authorisation within most Grids is performed using a Public Key Infrastructure, requiring each user to have an X.509 public-key certificate. Authentication is performed through ownership of a particular certificate, while authorization decisions are made based on the user's identity (derived from their X.509 certificate), membership of a particular VO, or additional information assigned to a user by a VO. While efficient and scalable, this approach has been found wanting in terms of usability – many users find certificates difficult to manage, for various reasons.

One approach to ameliorating this issue, adopted to some extent by Grid communities already, is to abstract away direct access to certificates from users, instead using alternative authentication mechanisms and then converting the credential provided by these into standard Grid certificates. Some implementations of this idea use existing federated authentication techniques. However, current implementations of this approach suffer from a number of problems, not the least of which is the inability

to use the federated credentials used to authenticate to a credential-conversion portal also to authenticate directly to non-web resources such as secure shell daemons.

The ability to use federated authentication directly, without the use of a credential conversion service, would allow users to authenticate a Grid and its associated services, allowing them to directly launch and control computing jobs, all without having to manage, or even see, an X.509 public-key certificate at any point in the process. Authorization within the Grid would still be performed using VO membership asserted issued by the user's identity provider through the federated transport.

5.3 High performance computing

High-performance computing (HPC) uses supercomputers and computer to solve complex computation problems most commonly associated with scientific research or computational science.

Access to HPC resources is typically managed through the use of user certificates. This requires HPC operators to issue certificates to users using a registration process that often duplicates identity management processes that already exist within user organisations. The HPC community would like to utilise federated identity to perform both the user registration and authentication functions required to use HPC resources, and so reduce by avoiding this duplication of effort.

The HPC community also have following additional requirements:

- Improved Business Continuity: In the event of operational issues at an HPC system at one organisation (for example, a power failure), users and jobs could be transparently moved to other HPC systems.
- Establish HPC-as-a-service: Many organisations who have invested in HPC systems want to make their systems easily available to external customers. Federated authentication facilitates this by enabling these customers to use their existing identity, user credentialing and support processes.
- Improve the user experience: Authentication to HPC systems is normally performed using user digital certificates, which some users find difficult to use. Federated authentication can provide better user experience by allowing the use of other types of credentials, without requiring technical modifications to the HPC to support these.

5.4 Web content & applications

The Web is now an essential and ubiquitous tool for delivering content and applications within, and between, organisations. Unsurprisingly the migration of content and applications to the Web has been accompanied by a requirement for managing access to these.

This requirement led to the development of technologies such as SAML, which has subsequently been widely adopted by communities known as federations for managing access to these web-based resources between organisations.

These web-based approaches have proved very successful in many respects, but the increasing scale of federations have highlighted some challenges. An example of such a challenge is identity provider discovery. A web content or application provider, when presented with a user requesting access to a restricted resource, must obtain information about this user to determine their authorisations in order to make an access management decision; the process of determining which organisation is able to make these claims is known as identity provider discovery.

In the general case, however, there is no mechanism that allows a provider transparently to identify an authoritative source of such information. Because the provider has no way of knowing which organisation the user is affiliated with, the user is usually presented with a list that enumerates the organisations that the provider will trust to issue access management information. Unfortunately this is not uncommonly a large number; in some federations, hundreds. With the advent of inter-federation this is likely to increase to many hundreds and possibly thousands.

5.5 Roaming network access

The requirement for roaming network access within the education and research community has been motivated by two factors:

- Massive adoption of mobile devices by users within the community, such as laptops and smartphones.
- Deployment of large-scale campus wireless LANs.

The intrinsically collaborative nature of research has meant that many users within the community have always been highly mobile. More recently, the growth of international study, the sharing of facilities such as classrooms and libraries, and online delivery of learning content within education have also increased the requirement for

‘anywhere, anytime’ access to the network. Finally, high bandwidth access to the Internet is also increasingly viewed as an essential utility, and organisations are keen to accommodate their visitors’ desire for this.

However, organisations are often connected to national research and education networks that wish to maintain a ‘private network’ status, owing to less burdensome regulatory requirements. This requires that such networks do not provide public access; access to the network must therefore be constrained to users from within the research and education community. This requires establishing that a visitor is affiliated with an authorised organisation. In addition organisations may also be compelled by regulation or policy to maintain accountability for use of the network; this requires some means of identifying users. One strategy to address these requirements is to require the visitor to identify themselves prior to accessing the network; however, this can impose a significant burden on an organisation’s IT support staff or helpdesks.

Many organisations would therefore greatly benefit from an access management system that enabled visitors to connect to the network automatically by proving their affiliation to an authorised organisation. In addition, users would prefer not to have to obtain different credentials for each organisation that they visit; in most cases, they would prefer to use a single credential issued by their own organisation.

5.6 Multi-domain network management and configuration

Internetworks are constructed from a number of networks operated by distinct administrative domains. An example of such an internetwork is the GÉANT service area that interconnects the European national research and education networks to each other, and to other international peer networks. These networks have for many years supported higher priority services, in addition to best effort forwarding.

Owing to a lack of standardisation, these services often require manual configuration of network elements across multiple domains. This imposes longer lead-times and reducing opportunities for customisation. There is increasing demand for more dynamic and customisable services that would enable their customers more easily and conveniently to obtain the network service needed for a particular requirement.

Developing these services requires that network operators are able to request, on a dynamic basis:

- Modifications to the configuration of network elements, such as routers, that are managed by other operators.

- Access to information describing the status (such as connectivity, utilisation, etc.) of network elements, such as routers, that are managed by other operators.

These multi-domain networks are often highly co-operative, but nonetheless network operators want to place appropriate constraints on the access management privileges accorded to, for example, the administrators and customers of other networks with whom there is no direct relationship.

6. Conceptual model

This section describes a conceptual model for organising an ABFAB infrastructure for education and research. This model is depicted in Figure 3 below.

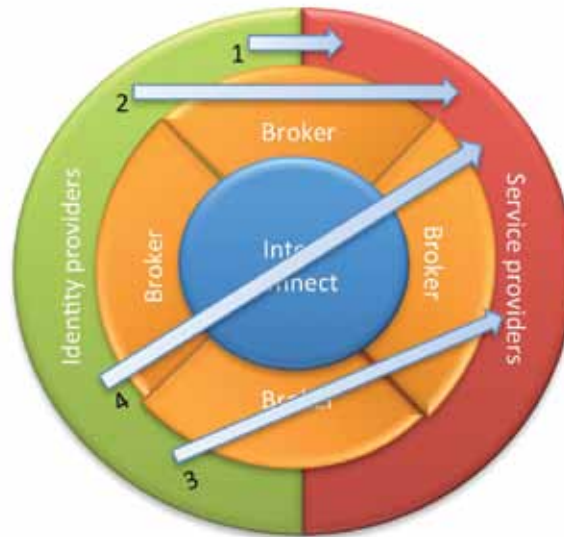


Figure 3

6.1 Identity Providers

Identity providers are organisations that maintain information about the identity of users and systems. Identity providers will typically be universities, schools and other organisations with affiliated users.

6.2 Service Providers

Service providers are organisations that consume information issued by identity providers. Service providers may often also be identity providers, but not always. For example, a cloud service provider is unlikely to be an identity provider; a university is likely to be both.

6.3 Brokers

Identity and service providers may often choose to connect to each other directly (arrow 1 in Figure 3). Frequently, however, it is more convenient and efficient to use

a *broker* that mediates between larger numbers of service and identity providers (arrow 2). Brokers create networks of connected identity and service providers.

Brokers themselves may also connect to each other, creating larger networks that include their respective identity and/or service providers (arrow 3). Other parties, such as communities of identity and/or service providers (e.g., an HPC consortium), may also accredit brokers to certify identity and/or service providers according to particular profiles (e.g. policies prescribing levels of assurance, eligibility and so forth).

6.4 Interconnect

Finally, a broker may act as a broker for other brokers. The union of these “broker brokers” is called the *interconnect*. Brokers within the interconnect are certified to a baseline profile that establishes a basic level of trust within the system, facilitating connectivity across the system (arrow 4) between otherwise isolated brokers.

7. Proposed implementation

This section outlines the work required to implement the organisation described in the previous section. This outline is primarily intended to provoke and inform discussion; it is likely that there are other strategies, and the community is encouraged to propose them.

The ABFAB architecture uses AAA technology (such as RADIUS or Diameter) as the basis of the connection infrastructure between identity and service providers. A widely deployed RADIUS infrastructure already exists. This infrastructure currently supports the global eduroam service but could be used to support ABFAB without significant modification (see section 11 for more information).

7.1 Interconnect

The international RADIUS infrastructure is composed of RADIUS systems that are typically nationally organised, with the national research and education network usually acting as the National Roaming Operator. This role is very similar to that of a broker, using the nomenclature of the previous section. The interconnect, in this case, is the set of regional top-level RADIUS proxy servers that connects the national RADIUS systems.

Although this infrastructure is essentially ready for an ABFAB-based system, some effort is nonetheless required to place the system on solid foundations. This work consists of the following:

- Establish a governance body.
 - *How:* Using the Global eduroam Governance Committee (GeGC) and its policies as a starting point, create a clean separation between the governance of eduroam (as an application of the interconnect) and the governance of the RADIUS infrastructure that is the technical manifestation of the interconnect. It is recommended that this new body is organised within a capable umbrella organization, such as TERENA or IGTF.
 - *Who:* Community consensus.
 - *Due:* November 2011.
- Develop an interconnect policy.

- *How:* Building on the existing eduroam policy, develop an interconnect policy that addresses both network and application access management.
- *Who:* The governance body, brokers, and other interested communities and stakeholders.
- *Due:* April 2012.

7.2 Brokers

The brokers, at least initially, would probably be largely represented by eduroam's National Roaming Operators (NRO) which are typically, but not exclusively, the national research and education networks. However, other bodies could assume the role of broker for non-eduroam applications for a particular constituency.

The interconnect's policy may impose modifications to the policies and practices of brokers

- Consider, and react to, the implications of the interconnect policy.
 - *How:* analyse the impact on existing policies, and potentially modify them. In particular, there may be significant efficiency savings to be made by re-using (insofar as possible) an existing national identity federation's framework to avoid duplication. SWAMID, the Swedish Higher Education identity federation, provides an excellent model for this that brokers are encouraged to consider. A substantial majority of education and research identity federations are also NROs, and so this approach would deliver a synthesis of policy, satisfying both network and applications access management, for their customers.
 - *Who:* any broker that wishes to participate.
 - *Due:* prior to any operational changes.
- Make any modifications to existing technical and business processes where necessary.
 - *How:* determined by results of the policy analysis.
 - *Who:* the broker, probably in co-ordination with their customers and user communities.
 - *Due:* prior to connection to the interconnect.

7.3 Identity and service providers

Identity Providers will need to:

- Consider, and react to, the implications of their broker's policies.
 - *How*: analyse the impact on existing policies, and potentially modify them.
 - *Who*: all providers, probably with the guidance of their broker.
 - *Due*: prior to any changes to operational practices.
- Make any modifications to existing operational practices where necessary.
 - *How*: determined by results of the policy analysis.
 - *Who*: all providers that intend to participate.
 - *Due*: prior to connection to their broker.
- Connect to a broker.
 - *How*: deploy ABFAB-conformant software (see Appendix III for more information).
 - *Who*: any provider that intends to participate, probably with the guidance of their broker.
 - *Due*: prior to any operational use.

8. Recommendations & roadmap

This paper proposes the following recommendations. These are principally directed at those stakeholders with a responsibility for organising and delivering access management services, and those that consume them.

1. The research and education community should seek to establish a common global access management system, based on the ABFAB technology, and develop a roadmap to achieve this goal. An outline roadmap is proposed in section 8.1 below.
2. Those stakeholders with a direct interest in access management should, given the strong possibility that the technology will be deployed within their constituencies in the near future, actively consider a strategy for managing this.

8.1 Roadmap

This paper proposes the following roadmap for the following six months, concluding with a workshop in October 2011.

1. Solicit feedback to this proposal.
2. Establish a mailing list to obtain feedback, and for general discussion of the proposal.
3. Create a wiki to collect and structure the feedback.
4. Identify alternative implementation strategies (to that described in section 7).
5. Develop a consensus strategy, based on this feedback.
6. Organise a workshop to seek consensus on an implementation strategy.

There is no Internet presence for this proposal at present; the manner in which this proposal is taken forwards – if at all – is a matter for the community. However, a temporary mailing list has been created to support co-ordination in the interim. Parties who are interested in tracking and discussing this proposal are encouraged to subscribe to cgams-announce@jiscmail.ac.uk at <http://www.jiscmail.ac.uk>.

9. Appendix I – ABFAB overview

This appendix describes the ABFAB architecture and related technology. As this section only provides an overview, more technically-oriented readers are strongly encouraged to read the ABFAB Architecture document.

The ABFAB architecture is therefore best understood as a novel arrangement of some tried-and-tested technologies. These technologies are:

- **Extensible Authentication Protocol (EAP).** EAP is an authentication protocol. The protocol is highly extensible; several dozen authentication methods have been defined for EAP and implemented and these support many different types of credentials (password, one time passwords, certificates, biometrics, SIM cards, etc). EAP is implemented within almost all desktop and mobile operating systems, although typically for network authentication only. EAP is already widely used within the education and research community for WiFi authentication in eduroam. *EAP is used within the ABFAB architecture to provide authentication of principals.*
- **Generic Security Services (GSS) API.** The GSS API provides a programmatic interface to a set of security-related functions, such as authentication, data integrity validation and confidentiality. The GSS API is widely used by applications, particularly for authentication-related services. Microsoft provides a very similar interface, called the Security Support Provider Interface (SSPI), that is widely used by Microsoft and other vendors for Windows-based applications. The Globus toolkit, which is widely used for Grid-based applications, makes extensive use of the GSS API. *The GSS API is used within the ABFAB architecture to provide application developers with an accessible strategy for integrating their applications into an ABFAB system. Specifically, the ABFAB architecture describes a GSS authentication mechanism that uses EAP for authentication and keying. The SASL GS2 mechanism also provides an integration strategy for SASL-based applications.*
- **RADIUS.** The RADIUS protocol is a long-established technology that enables authentication and authorisation functions to be de-coupled from a service, and logically moved to a central authentication, authorisation and accounting (AAA) server. The service, therefore, does not need to maintain credentials or policies; instead, these can be managed centrally, which may be a significant convenience within systems composed of many services. The RADIUS protocol also supports the ability for an AAA server to proxy protocol messages to other AAA servers –

possibly operated by organisations – thereby providing federation. This federation capability of RADIUS is used to great effect by eduroam, which performs many hundreds of thousands of international federated authentications per month. *The RADIUS protocol is used within the ABFAB architecture to federate EAP authentication.*

- **Security Assertion Mark-Up Language (SAML).** SAML is a framework for encoding and exchanging security-related information. This information typically describes authentication events and any authorisations associated with the user being authenticated. SAML has become widely deployed to support Web-based applications. *SAML is used primarily within the ABFAB architecture to encode authorisation information.*

Figure 4 below illustrates the composition of these technologies.



Figure 4

The Client and Authentication server (or AAA server) is part of the same organisation (as indicated by the light blue tint). The Client has an Identity Selector, which is software that manages the user's identity – or identities if, for example, the user is affiliated with more than one organisation.

An identity has some credentials associated with it; for example, a user-name and password, or a certificate. A client application, if required by a service application to authenticate, obtains some credentials for the user through the Identity Selector (depending on policy, the user may or may not be prompted to select an identity to use).

These credentials are authenticated (red arrow in Figure 4) against the user's home EAP server using an EAP authentication method. The user's EAP credentials

are securely protected between the Identity Selector and the EAP server during authentication, typically using TLS. The EAP credentials are transported within the application protocol (e.g. SSH or HTTP) between the Client and Server applications; and within RADIUS between the Service and EAP server.

If authentication is successful, the process concludes with the Client and Service having possession of the EAP Master Shared Key (MSK). This key is used to establish a GSS security context between the Client and Server applications. If SAML information is provided during EAP authentication, this is made available to the Server application within the GSS context.

10. Appendix II – Technical feasibility

Project Moonshot is presently undertaking an implementation of ABFAB. The implementation work can be divided into two parts: infrastructure and applications.

10.1 Infrastructure software

The infrastructure software is primarily concerned with concealing the details of the architecture from applications and users. This is achieved by means the following components:

- The **GSS EAP library** provides a complete GSS EAP mechanism for both the MIT and Heimdal GSS stacks. The mechanism has been built and tested successfully on Linux, Mac OS X and FreeBSD. A Windows port is planned and will be completed by Q3 2011.
- A **Secure RADIUS (RadSec) library**, which has been derived from radsecproxy. This library is used by the GSS EAP library for RADIUS support. The library is mainly complete; some additional work is required to improve stability and resilience.
- The **Shibboleth Service Provider**, which is used by the GSS EAP library for SAML processing. This enables the GSS EAP library to expose SAML constructs to applications using GSS naming extensions.
- The **Moonshot Identity Selector**, which enables a user to manage his identities and select one, if required by an application. The Identity Selector is currently being implemented for Linux and Windows.

This experience indicates that, while the infrastructure code is relatively complex to implement, the ABFAB architecture builds easily on, and integrates cleanly with, the operating systems' existing security software. The infrastructure software is currently being packaged for Debian Linux.

10.2 Application software

The applications part concerns software applications that wish to enjoy the benefits conferred by the ABFAB technology. The following applications have been successfully tested:

- **Apache.** A new GSS authentication module has been implemented, derived from *mod_auth_kerb*. The new code will be merged back with *mod_auth_kerb* by Q3 2011, which many Linux distributions already ship.
- **Firefox.** A new GSS authentication plug-in has been implemented, derived from Firefox's existing support for the HTTP Negotiate authentication scheme. The new code will be merged with Firefox's existing Negotiate support by Q3 2011.
- **OpenSSH.** No client modifications required; minor server modifications required.
- **MyProxy.** Very minor modification required.
- **OpenLDAP.** No client or server modifications required.
- **Adium.** No modifications required.
- **Jabberd.** No modifications required.

This experience suggests that many, or even most, applications will require little or no modification to work with the ABFAB architecture. Only the most complex applications, such as Apache and Firefox, have required non-trivial work.

10.3 Utility

Early results strongly indicate that the ABFAB technology is capable of addressing the use-cases described in section 4.

10.3.1 Cloud

A significant number of services could potentially be delivered from Cloud-based infrastructure. It is, therefore, necessary to be circumspect when making generalisations about the applicability of ABFAB in this context.

However, on the basis of the evidence described in section 10, it seems likely that the majority of applications supporting either GSS-API or SASL will require little or no modification. Although the existing sample size of applications is small, it is worth noting that the applications requiring the least (or no) modifications were SASL-based applications; and it is these applications – in particular, mail, directories and messaging – that are typically cited as being of most interest for Cloud.

10.3.2 Grid infrastructure

As described previously, ABFAB-based authentication has been demonstrated with appropriately modified OpenSSH (server modifications only) and MyProxy. SAML-based authorisation was also demonstrated with OpenSSH. Using the modified MyProxy it was possible to obtain an X.509 certificate for the authenticated identity, and also retrieve a proxy certificate from the repository.

10.3.3 High performance computing

As described previously, ABFAB-based authentication and authorisation has been demonstrated with appropriately modified OpenSSH (server modifications only).

10.3.4 Web content & applications

As described previously, ABFAB-based authentication has been successfully demonstrated using appropriately modified versions of Firefox and Apache. The identity provider discovery problem is substantially addressed through the use of the client-based identity selector: in the worst case, the user is only presented with a list of those identities that he already possesses, rather than all of those identity providers trusted by the service provider. These modifications will be merged back during Q3 2011 so that the functionality will be widely available. Some further work to the Shibboleth SP, planned for Q3/Q4 2011 by Project Moonshot, is necessary to achieve equivalence between ABFAB and conventional Web SSO authorisation.

10.3.5 Roaming network access

The success of the eduroam service convincingly demonstrates the applicability of RADIUS-based access management to roaming network access. As proposed in section 7, an ABFAB-based system could potentially re-use the international RADIUS infrastructure that currently supports eduroam: enabling a common infrastructure for application and network access management. Some of the specific ABFAB-related technology could potentially be extended to eduroam; for example, SAML attributes provided by a RADIUS server could potentially be used to provide authorisation capabilities.

10.3.6 Multi-domain network management and configuration

As described previously, ABFAB-based authentication and authorisation has been demonstrated with appropriately modified OpenSSH (server modifications only).

11. Appendix III – Expected modifications to RADIUS infrastructure

This proposal suggests using the international RADIUS infrastructure that presently supports eduroam as the basis for this system. Some minor modifications to the system will be necessary; this appendix briefly describes these.

The existing infrastructure does not require enforcement of name constraints on service providers; which is to say, a service provider can today claim any name. In the particular case of eduroam, this deficiency does not present any problems because authorisation in eduroam is not contingent on the name of the wireless access point; indeed, wireless access points do not generally have globally meaningful names. Consequently, the EAP peer is able to infer, from a successful authentication, that the wireless access point is authorised to provide an eduroam service; but not that it is a particular wireless access point. In the case of eduroam, the latter distinction is not necessary: an access point is providing an equivalent service to any other access point.

In an ABFAB system it is essential, in the general case, that the EAP peer is able to disambiguate between different services. For example a user, connecting to an SSH daemon on a server, almost certainly requires assurance that it is connecting to a particular server, and not any server.

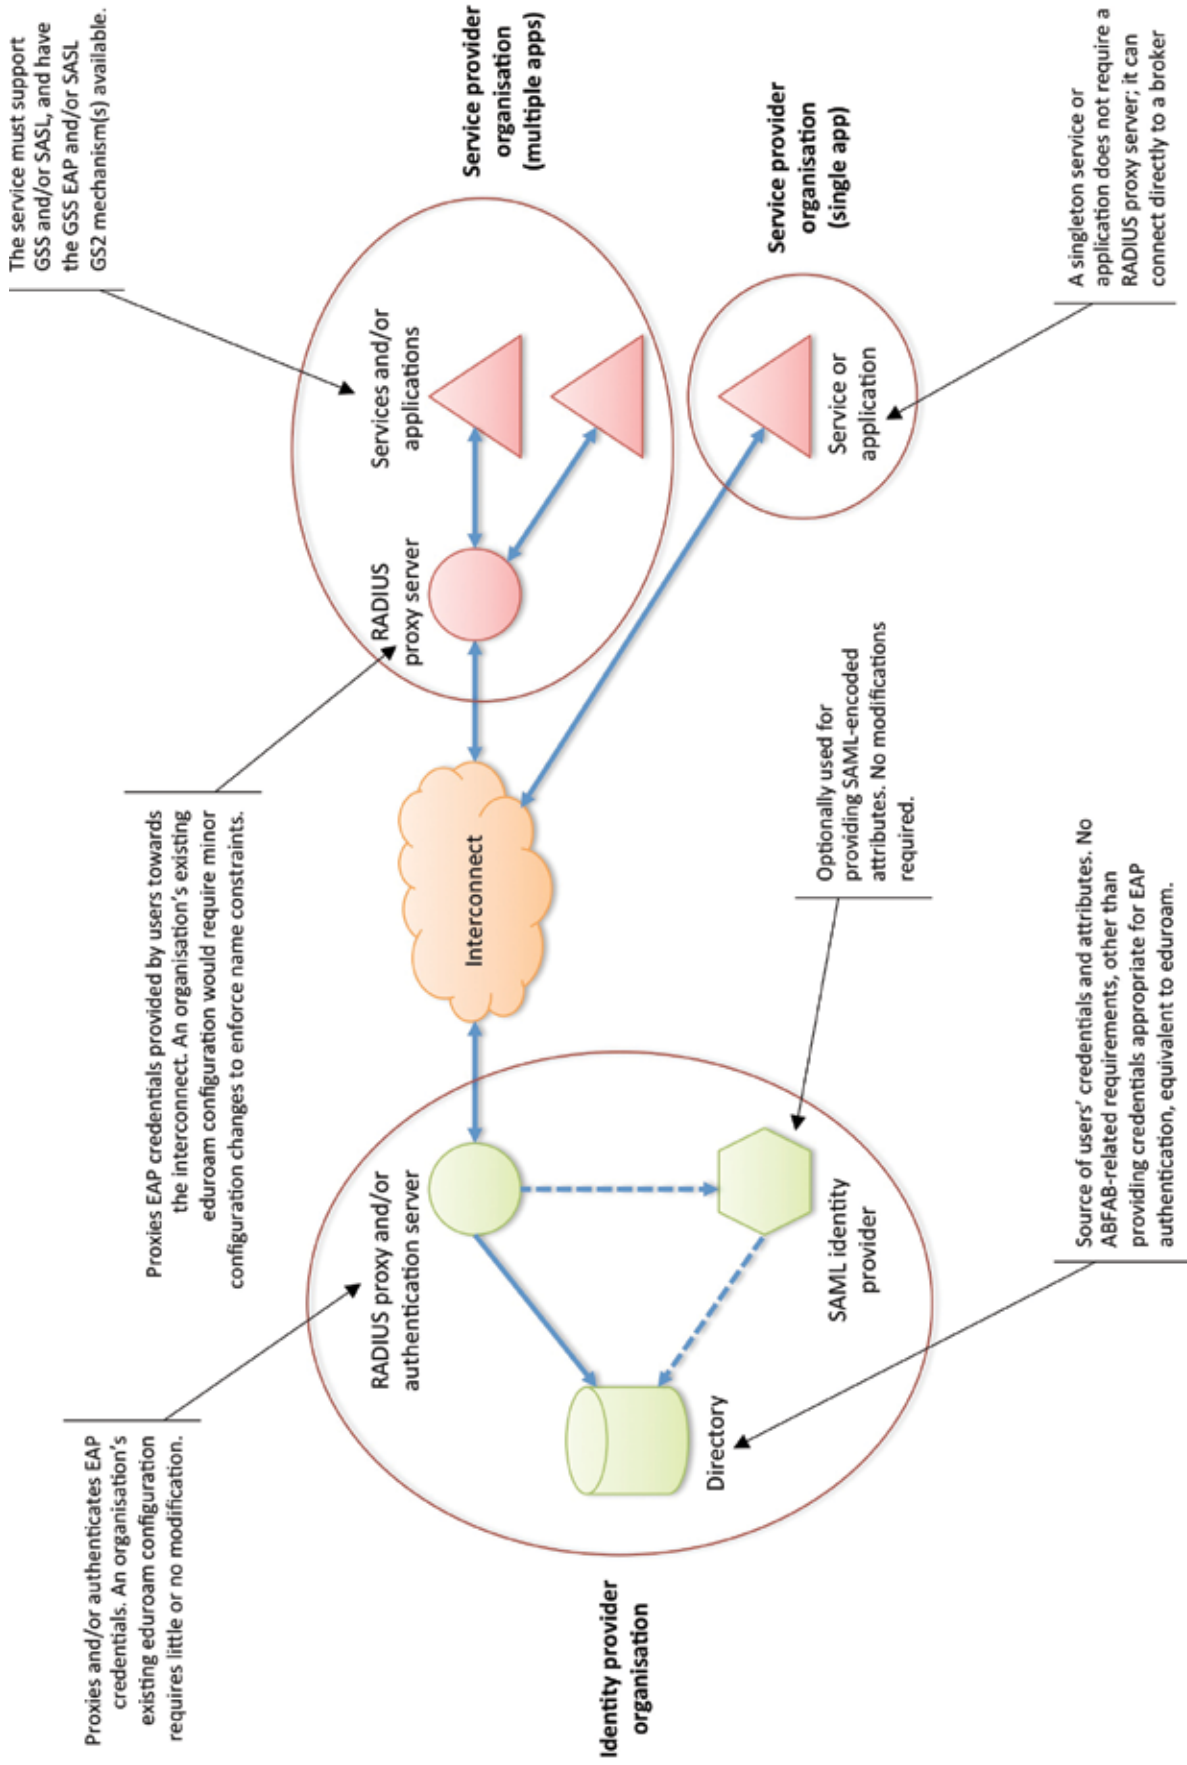
Consequently, ABFAB requires the use of EAP channel bindings. This provides a mechanism that enables the EAP peer to signal, within the EAP authentication method, to the EAP server which service it believes it is connecting to. If the EAP server can match the EAP peer's opinion of the service's name against the RADIUS system's same opinion, then it concludes that the service is authorised to wield this name. If there is a mismatch, then the EAP server rejects the authentication. EAP channel bindings will be implemented in the Moonshot Identity Selector (the EAP peer actor) and FreeRADIUS (the EAP server) by the end of Q2 2011.

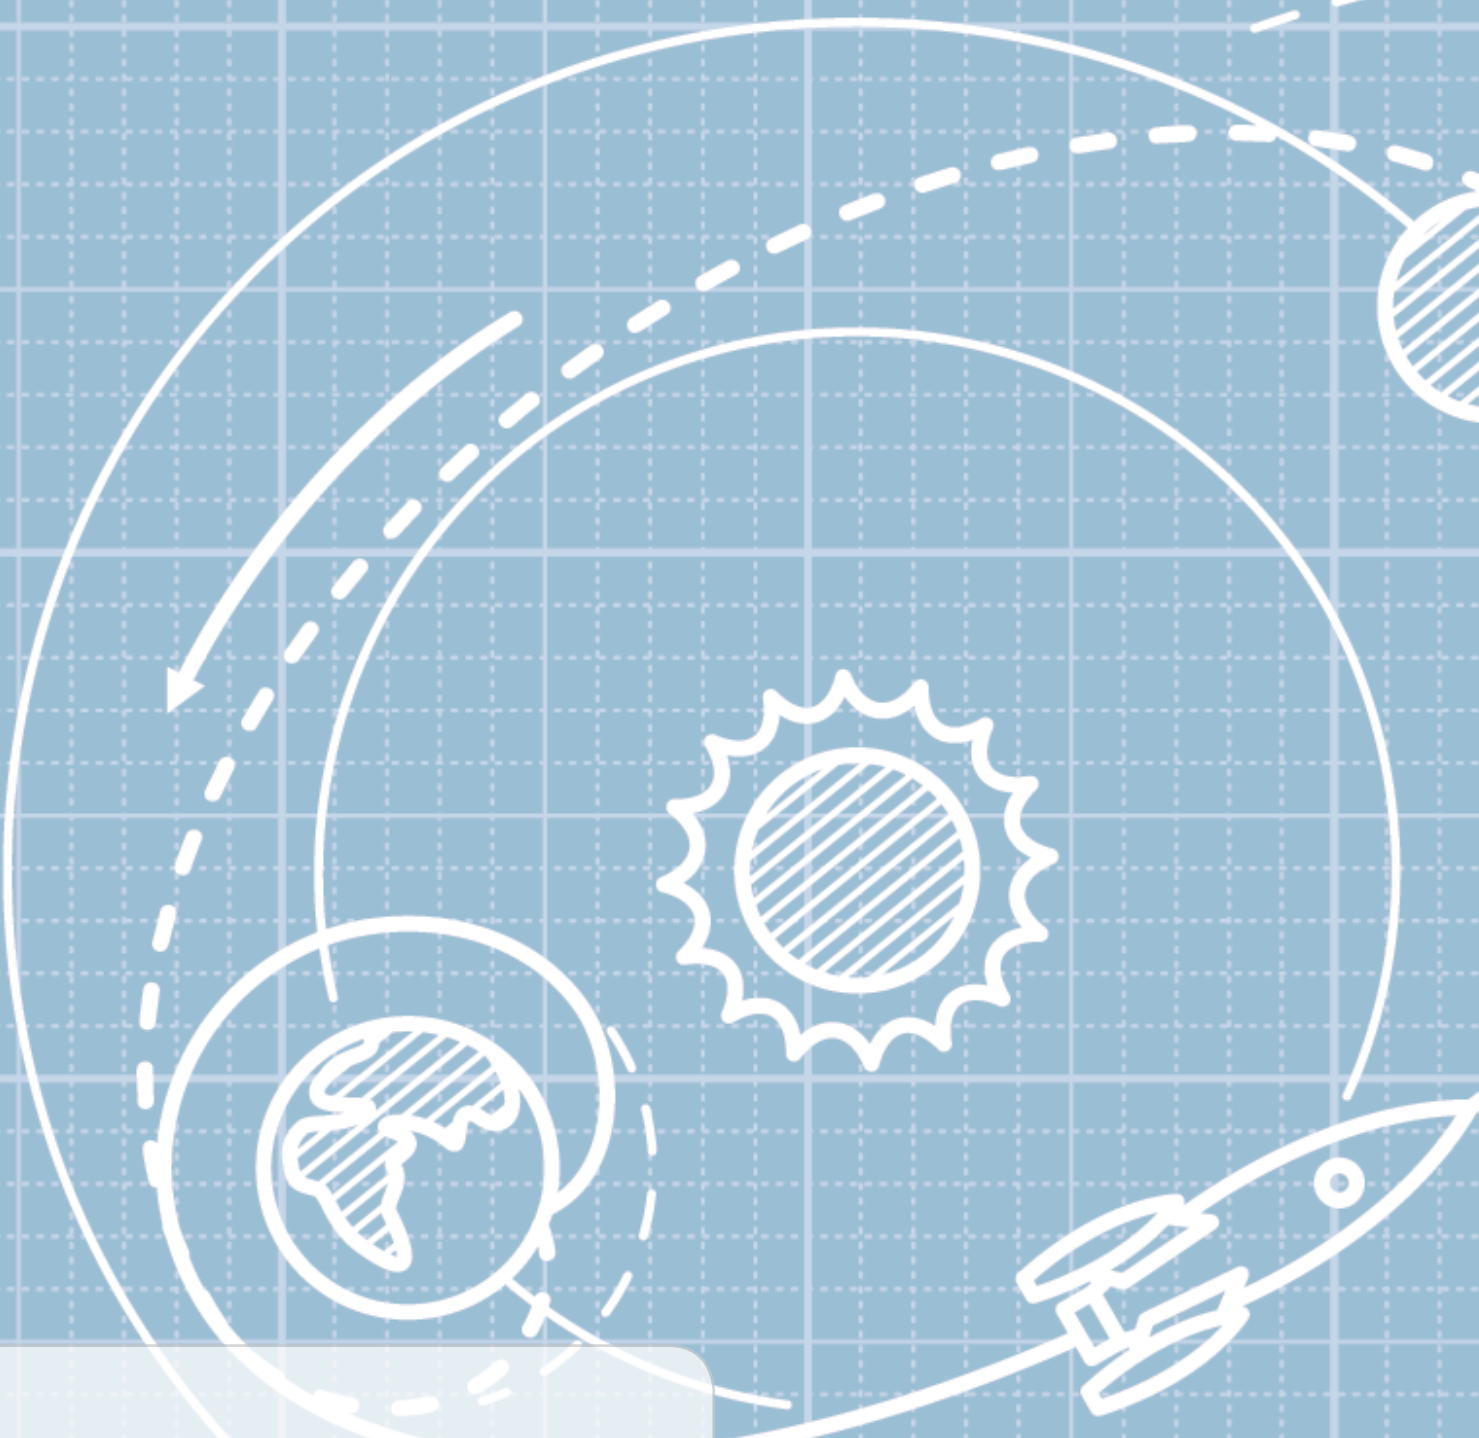
Therefore, identity providers will only be required to upgrade their RADIUS server software.

The EAP server obtains the RADIUS system's opinion of the service's name from equivalent channel binding data transported in the RADIUS layer. For this data to be useful for the bindings comparison, an actor trusted by the EAP server must validate it. In this case, the actor will be the broker that registered the service provider.

Therefore, brokers would be required to validate the channel binding data (expressed in RADIUS attributes) claimed by service providers. This is simply a matter of matching the claimed name against the service provider's registered name(s). However, this information must also be collected and maintained. This should not, however, be particularly onerous: eduroam's NROs already validate and maintain the names associated with identity providers, and so no further action is required for service providers that are already identity providers. In the case of most if not all NROs this information would need to be obtained for the remaining service providers that are not identity providers.

Organisations hosting multiple services behind a single proxy would be required to validate the channel binding data in the same way. Similarly, it is expected that the collection and management of this information would be a trivial exercise in most cases.





Parties who are interested in tracking and discussing this proposal are encouraged to subscribe to cgams-announce@jiscmail.ac.uk at <http://www.jiscmail.ac.uk>.

This document is copyright the JNT Association 2011.

Published by The JNT Association,
Lumen House, Library Avenue,
Harwell Oxford, Didcot,
Oxon OX11 0SG
United Kingdom
www.ja.net

$$\frac{M_1}{M_2} = e$$

$$V =$$