

Federated Global Identity Management: Towards a Framework

Jawed Siddiqi, Babak Akhgar and Mehrdad Naderi: Sheffield Hallam University (UK)

Wolfgang Orth: Fraunhofer Institute for Secure Information Technology (Germany)

Norbert Meyer: Poznan Supercomputing & Networking Center (Poland)

Miika Tuisku: University of Helsinki (Finland)

Gregor Pipan: XLAB (Slovenia)

Name of Conference: **The 2006 International Conference on Grid Computing & Applications (GCA'06)**

Abstract - The emerging debate over identification and the selection of technology to authenticate individuals and enterprises alike is among the most important issues that is shaping the information age today. A review of the current state of the art of research around digital identification on both the Web and the Grid is presented as a background in order to propose a novel and innovative framework for federated global identity management. The paper presents a progress report on the FEGIMA project that aims to deliver a unique framework whose fundamental innovation lies in the fact that it extends the management of global identity to get federated access across heterogeneous ICT platforms.

Keywords: **FEDERATED IDENTITY ARCHITECTURE AND FRAMEWORK**

1. Introduction

The emerging debate over identification and the selection of technology to authenticate individuals and enterprises alike is among the most important issues that is shaping the information age today. Clearly, the concept of identity is far broader than the mere content of a name. While names and naming protocols are a critical element of identity, in that they provide the means to distinguish one individual from another, the underlying relevance, role, context and meaning attributed to a given individual can only be gleaned by reference to other factors. In the human space, this is due to people existing in many social, economic, political, cultural and other dimensions concurrently. Whilst in the digital space the varying architectures (legacy and new) and underlying enabling technologies pose a similar dilemma. In defining the identity of a person be it in human or digital spaces which is termed as virtual identity is a multi-faceted complex problem.

Current identity and security management in ICT platforms is dedicated, restricted, and limited to individual users and applications. The FEGIMA project aims to deliver a unique framework whose fundamental innovation lies in the fact that it extends the management of global identity to get federated access across heterogeneous ICT platforms. Further innovation in the FEGIMA framework is the integration of this with trust and policy management that is reinforced by an intrusion detection system. Here we primarily report on the framework for federated global identity management.

The research was led through a collaborative effort by the following four partners: Sheffield Hallam

University (UK), Fraunhofer Institute for Secure Information Technology (Germany), University of Helsinki (Finland) and Poznan Supercomputing & Networking Center (Poland) leading to the proposed development of the proposed FEGIMA framework. This was then distributed for peer review and comments to our industrial partners that included Sun Microsystems (Belgium), Telecom Italia (Italy), Telefónica Investigación y Desarrollo (Spain) Valimo Wireless (Finland) and an SME XLAB (Slovenia). Therefore, the proposal benefits from being a true European wide collaborative effort in development and equally importantly a peer to peer critique between leading academic institutions and key industrial partners involved in security and telecommunications.

In the next section a state of the art research on digital identification, related to both the Web & Grid global infrastructures. In the former (ie the Web) where there is considerable volume of literature we have selectively chosen two governing bodies of Liberty Alliance and Shibboleth are reviewed, whilst in the latter (ie for the Grid), where there is relatively sparse volume of literature, the issues of authentication and authorisation are considered. The remainder of the paper provides an overview of the proposed framework and the associated architecture. Finally, the proposal is then discussed in terms of challenges and innovation that are present in the FEGIMA project.

2. State-of-the-Art Research on Digital Identification

The following surveys the relevant issues relating to identification or identity management systems both in the public and private sectors. One key similarity is

that both the public and private sectors wish to enable a system that will allow an end user (whether an individual or organisation) to enjoy the convenience of “single sign-on.” (SSO) whilst increasing the opportunity to discover fraud against the systems. It remains to be seen whether there are systems and processes that can be used across both sectors. While current architectures appear to be primarily or exclusively suitable in one or the other sector, it is clear that ultimately there will be a sufficient demand for cross-sector interoperability that common-denominator solutions will be required.

Presently federated identity within one global computing infrastructure e.g. web space continues to be a significant challenge that is the focus of sizeable volume of research. However, in contrast the research into federated identity in the grid space is relatively sparse.

2.1 Electronic Identity Management on the Web

In this section we review key aspects of work, relevant to FeGIMa, for two of the key governing bodies that are striving towards standardisation for identity management on the web, namely Liberty Alliance and Shibboleth.

Liberty Alliance:

The Liberty Alliance [1] is a consortium of more than 150 organisations that develops specifications for federated identity management. It is working on the development, deployment and evolution of an open, interoperable standard for network identity where privacy, security and trust are maintained.

Liberty Alliance contributions towards federated digital identity management has been the development of three specifications:

- Identity Federation Framework (ID-FF); enables identity federation and management through features such as identity/account linkage, simplified sign on and simple session management
- Identity Web Services Framework (ID-WSF); provides the framework for building interoperable identity services, permission based attribute sharing, identity service description and discovery and the associated security profiles
- Identity Services Interface Specifications (ID-SIS) enables interoperable identity services such as identity profile services, alert services, calendar services, wallet services, contacts services, etc.

Federated identity and attribute servers are the solution used by Internet2 [2] projects and were developed for individual virtual organizations especially within the European Data Grid (EDG) [3]. Within federated identity management authentication depends on the notion of identity. The Liberty Alliance project's

Identity Federation Framework ID-FF specifies a third party authentication model, where individual services rely upon assertions (SOAP messages carried over HTTP) generated by an identity provider. Thus, the service is not required to directly authenticate the user, but rather an entity whose sole responsibility is to identify the user based on direct authentication. This model thus requires that the service provider trusts the identity provider.

The Liberty Authentication Service of the Web as an exemplar allows SOAP client applications to authenticate to the service via any of the authentication models specified by the IETF's Simple Authentication and Security Layer specification, thus standardising the authentication methods used.

Liberty explicitly accommodates identity provider use of arbitrary authentication mechanisms and technologies. Different identity providers will choose different technologies, follow different processes, and be bound by different legal obligations with respect to how they authenticate users. The choices that an identity provider makes here will be driven in large part by the requirements of the service providers with which the identity provider has federated.

Within the context of Liberty Alliance Web Service Framework, for identity providers and service providers to communicate with each other, they must a priori have obtained metadata regarding each other. These provider metadata include items such as X.509 certificates and service endpoints.

When sharing a globally known identifier among separate organisations, the users privacy may be compromised. Liberty allows the creation of opaque privacy-protected name identifiers, which may cross organisations without compromising the privacy of the user or leaking data.

Shibboleth:

Shibboleth [4] is a joint project of Internet2/MACE (Middleware Architecture Committee for Education) and IBM. Its focus is to investigate architectures, frameworks, and practical technologies to support inter-institutional sharing and controlled access to web available services. Shibboleth focuses on inter-institutional resource sharing within academia, but the project is relevant to many business settings as well. The project will produce an analysis of the architectural issues involved in providing such inter-institutional services, given current campus realities and the current state of relevant standards. It will also produce a pilot implementation to demonstrate these concepts.

Shibboleth utilises frameworks for multiple, scaleable trust and policy sets, termed as Clubs, to specify a set of parties who have agreed to a common set of policies. This moves the trust framework beyond bi-lateral agreements, while providing flexibility when different situations require different policy sets. It does

so by making use of open SAML [5] for the message and assertion formats, and protocol bindings.

Key concepts within Shibboleth include:

- **Federated Administration:** The originating campus provides attribute assertions about the user to the target site. A trust fabric exists between these, allowing each site to identify the other, and assign a trust level. Originating sites are responsible for authenticating their users, but can use any reliable means to achieve this.
- **Access Control Based On Attributes:** Access control decisions are made using those assertions. The collection of assertions might include identities, however, many situations will not require this (eg accessing a resource licensed for use by all active members of the campus community or accessing a resource available to students in a particular course).
- **Active Management of Privacy:** The original site and the user, control what information is released to the target. A typical default is merely "member of community". Individuals can manage attribute release via a web-based user interface and are no longer at the mercy of the target's privacy policy.

2.2 Electronic Identity Management in the Grid

Grid computing has emerged as an important new field, distinguished from conventional distributed computing by its focus on flexible, secure, co-ordinated resource sharing among dynamic collections of individuals, organisations and resources. This sharing is, necessarily highly controlled, with resource providers and users defining clearly and carefully what is shared, who is allowed to share and the conditions under which this sharing occurs. Ensuring security in such settings can be categorised into several fairly independent areas of identity management: authentication; authorisation; secure communication; auditing and accountability and intrusion detection. For the purposes of the framework presented here we focus on the first two areas namely: authentication and authorisation.

Authentication

Today emerging grid security efforts are beginning to address application and infrastructure security issues including application protection and node-to node communications. Among other advances, emerging grid security approaches are integrating Kerberos security with PKI/X.509 mechanisms, securing peer connections between network nodes and better protecting grid users and applications from malicious or badly formed code.

Thus far Kerberos [6] has been an early and still viable solution to global identity and authentication in somewhat restricted environments, while PKI [7] identities combined with the Transport Layer Security

(TLS) [8] protocol is the solution widely adopted in the Grid.

Grid Security Infrastructure (GSI) in the form of Globus as well as Unicore's security model utilise public key cryptography, specifically public/private keys and X.509 certificates as the basis for creating secure grids; however the authentication mechanisms of both systems differ. While UNICORE signs each part of the job with the user's certificate, which guarantees the integrity of jobs and authenticates the submitting user of a job and therefore enables end-to-end security model, Globus toolkit, uses long-term X.509 certificates to generate a temporary proxy that can act on a user's behalf without requiring user intervention. Once created, the proxy is used to grant or deny access to resources found throughout the grid thus enabling delegation. Because the proxy is used across system, this gives the end user the ability to sign on only once. The proxy expires within a preset amount of time.

In order to provide interoperability between UNICORE and Globus solutions, the GRIP project addressed the following key aspects: translating UNICORE requests for job submission, output retrieval, and status queries to the corresponding Globus constructs and mapping of permanent UNICORE user certificates to temporary Globus proxy certificates.

Within this project these functions were to be implemented without changes to the respective architectures. The result of the project is the development of the Enhanced Target System Interface (ETSI) that enables submitting jobs from Unicore clients to Globus systems and return the results of the computation to the users.

Globus and Unicore are not directly compatible with Kerberos authentication. One emerging grid security effort attempts to sidestep this problem by blending Kerberos infrastructure and X.509 certification. The so-called KX.509, developed at the University of Michigan, is designed to provide a bridge between Kerberos and PKI.

Therefore one can see that the use of a common format credentials with Grid middleware implementations is not yet seamless. Even though majority of the heterogeneous grid middleware implementations are based on the Globus toolkit and most of these implementations utilise X509 certificate extensions to address security and access to the virtual grid resources in the respective grid middleware implementations. Considering EDG and Globus as an exemplar situation, neither of these are interoperable with the other. The best situation today is that the new versions of the middleware ignore the certificate extensions that they do not recognise. Rather there are many different solutions to different parts of the problem. One of the reasons for so many solutions is that authentication is the first critical step to any trusted use of resources. Grid authentication must also interact with user and grid resource authentication requirements; therefore, a single monolithic solution is not feasible. The

paradigm of federating authentication from various servers and mapping credentials between a common Grid middleware is the most promising solution as proposed by the FeGIMa project.

Authorisation

Several research collaboration efforts have addressed the authorisation requirements of distributed computing and collaboration. The requirements vary drastically depending on the application. Authentication systems must be simple for users, access policy should be transparent to the user of resources, easy to set and maintain by the owners of individual resources and site administrators.

In Globus a user proxy requiring access to a resource first determines the identity of the resource proxy for that resource. It then issues a request to the appropriate resource proxy. It is up to a resource proxy to enforce any local authorization requirements. Depending on the nature of the resource and local policy, authorisation is checked and if the request is successful the resource is allocated and a process created on that resource.

The verification requires mapping the user's credentials into a local user id or account name. In a GSI enabled grid, the system receiving the request reads the user's name from the proxy, and then accesses a local file to map the name to a local user.

To avoid creating scores of extra user IDs on different grid systems, administrators can assign users to virtual groups. All users from a particular domain can be mapped to a single, common user ID when accessing a given grid resource. GSI is designed this way to help administrators separate outside users running grid computations from local users in need of local administration and support.

In Unicore security model, certificates serve as grid wide user identifiers, which are mapped to local account at each Unicore site. In addition the site retains full control over the acceptance of users based on the identity of the individual, the distinguished name or other information that might be contained in the certificate. Each site can restrict and limit accessible resources at each target systems, thus retaining the ultimate control.

One of the main authorisation issues is how to name users in a manner that is meaningful both at the resource site and across the Grid. The adoption of FeGIMa federated identity management will address the solution here. The research agenda and details of the process adopted for the FEGIMA project are detailed in a companion paper [], here we provide progress so far of the technical details that have been developed.

3. Overview of framework

In designing the FeGIMa framework for the Global ICT infrastructures of the Grid and the Web, the focus of research will be on three major building blocks (i.e. areas of work), those being:

- Federated Identity Management
- Trust Management and Policy Execution
- Intrusion Detection

In developing the FeGIMa architecture, underpinned from the above three FeGIMa framework building blocks are sub-systems that will be based on the following:

- **Base Technologies** will provide the necessary functionality which if missing would render the whole framework ineffective i.e. without the ability of securing communications the goal of privacy is seriously endangered, no matter how powerful the privacy policies, the user is able to define. However these technologies alone are not sufficient to provide an encompassing solution to the problem of federated identity management.
- **Core Components** contain the central research targets of the project that will be developed by the consortium partners. These Components are necessary to realise federated identity management in its full scope. They will be integrated with the base components to leverage their functionality, exemplar core components include queries regarding policy or privacy information should be digitally signed and encrypted.
- **High-Level Services** are services and applications that will be realised, when the base and core components are developed by the FeGIMa project. They will provide the stage to develop commercially exploitable use cases that will facilitate integrated federated global identity management in the Grid and Web spaces.

To realise the above vision, the project will initially develop the Federated Identity Management component for Grid environments. Subsequently, the integration of the prepared subsystem with already existing Web federated identity systems will be developed. Thus realising the global identity management system. Together with the Trust Management System and the Intrusion Detection System, will compose the integrated Federated Global Identity Management Framework. (see Fig 1)

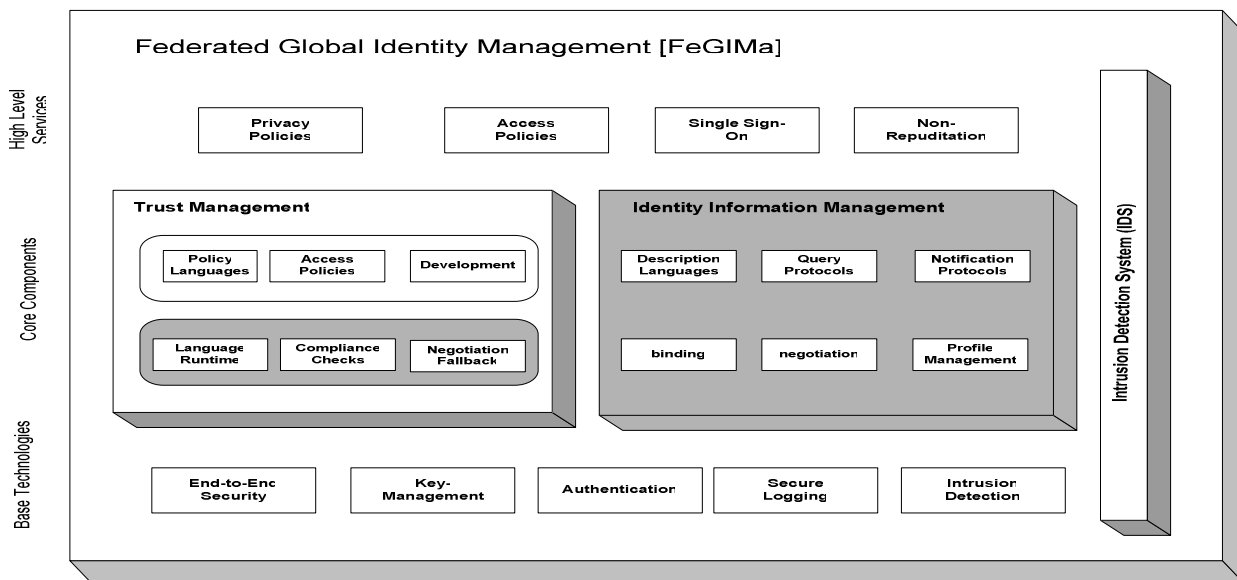


Figure 1: FeGIMa Framework

The realisation of this FeGIMa framework is illustrated in Figure 2: FeGIMa Architecture, illustrating the integration of hither to disparate global ICT infrastructures of the Grid and the Web in an open and adaptive framework that integrates these security

mechanisms without endangering the separated and independent grids or the web. FeGIMa is achieved through a comprehensive security management framework that is extensible and flexible and utilises open standards. The provision of identity information is handled by framework components that employ their respective protocols from the relevant standards allowing for a simplified access to Grid systems or the Web for users.

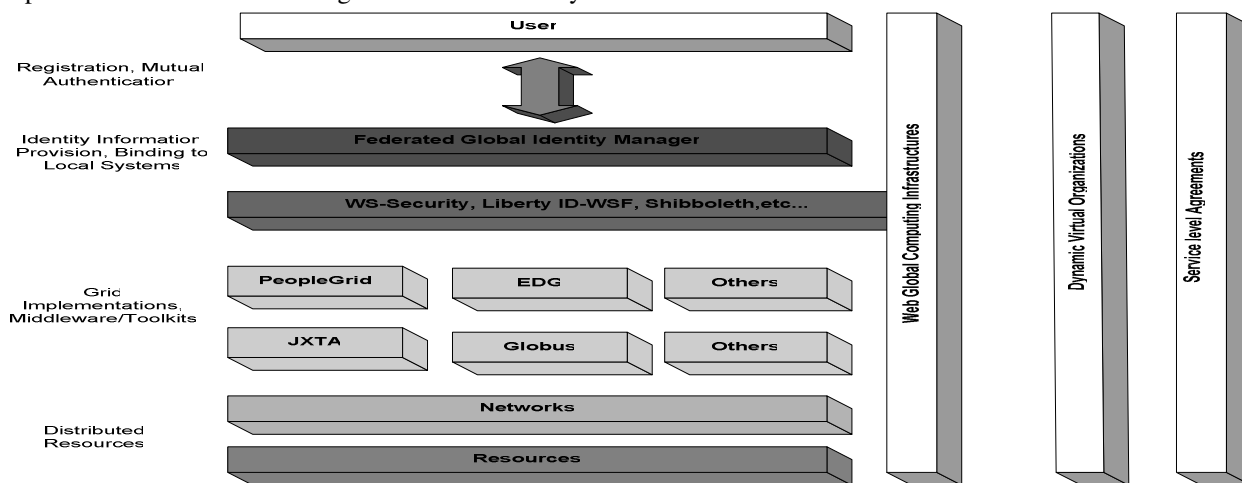


Figure 2: FeGIMa Architecture

Next is a selective discussion of the challenges and innovation in the project that focuses on federated identity management within the Global ICT infrastructures of the Grid and the Web.

4. Challenges and Innovation

To meet the challenge of federated global identity management through the paradigm of Single Sign On (SSO); thus, reducing the number of separate credentials clearly requires crucial components of strong authentication and credential management.

Beyond authentication the need exists to link applications and services, creating persistent and secure sessions to access these applications and services across global ICT infrastructures.

The current state of network identity in the Web requires the user to maintain individual islands of identity where the individual is responsible for remembering the multiple username/password pairs for each of these identity islands, and they must also manage the information that each site maintains in order to ensure that it is both up-to-date and appropriate. Within theGrid the same situation applies to a degree where each grid middleware implementation maintains its own security and identity management system and even where these middleware implementations are based on the same core systems they still do not recognise each other's security credentials. The best that has been achieved

as discussed in the state-of-art section of this proposal is the recognition of certificates in middleware that are based on a common implementation (i.e. Globus toolkit, EDG, and NorduGrid) where they ignore each other's certificate extensions or the development of rudimentary bridges (e.g. Globus, Unicore and Grip project).

The proposed FeGIMa Global Federated Identity Management Architecture will address these issues through removing from users the burden of maintaining their identities by providing the trusted hosting service (FeGIMa identity provider (FeGIMa IP)) that will manage the authentication and identification process on behalf of the user (identity owner). The proposed FeGIMa IP will manage authentication and identification process via a number of mechanisms providing weak to strong degree of trust based on the level of service required by the consumer (service provider) in the identification of the user .

The term 'federation' refers to the proposed FeGIMa framework that will make identity and entitlements portable across autonomous policy domains within the global ICT infrastructures of the Web and Grid, thereby overcoming the issues raised consequently federated identity is portable identity "across" these global ICT infrastructures and "within" the same infrastructure. Furthermore, the development of federated relationships between organisations means users (identity owners) have the freedom and the flexibility to move more seamlessly from one service provider (consumer) to another.

The existing approaches of providing security on the middleware layer (Grid) and application layer (Web) deliver intermediate solutions that are not suitable for addressing the needs of roaming mobile users from a user perspective. Furthermore, utilising functionality provided by higher layer services developed by the FeGIMa project will significantly reduce the complexity of interoperable security mechanisms required in the GC Grid infrastructure layer and avoid the co-existence of several independent infrastructures.

The vertical issues of FeGIMa will be based in the analysis of the overall approach with respect to end-to-end attributes, ranging from networking issues until the final orchestration of the services provided to the end-user. Business models that will be applied in the FeGIMa project, will affect all of the layers that are presented in the architecture. It is thus necessary to define all the aspect that are involved in this adaptable configuration for the end-to-end specific attributes of the FeGIMa architecture.

First services are defined on the application support layer providing also an interface towards the user, which is achieved either through direct interaction of offline by means of a contract with a service provider through a Service Level Agreement (SLA). Based on this, negotiation, control and management of the

SLA need support for service providers for example as part of a hosting environment and also for a client that needs maybe a service from a different provider that supervises the execution and SLA violations.

Next, the session management and user identity management, which is achieved by the network middleware layer and the introduction of SLA support which will itself require the transference of monitoring and performance information from the underlying networking layer to the Application layer. Clearly, the monitoring and support of the SLA cannot be based on the layered attributes of the current architecture, but it has to be based transparently on the vertical 'signalling' and session management between the subsystems in each layer.

In this context, the understanding a Virtual Organisation (VO) as an organisational unit bringing together partially the resources offered by different parties clearly poses challenging problems with respect to control, access and the availability of resources provided to the VO's. This issue will be addressed within FeGIMa by establishing an information exchange between the Administrative Domain organisational unit from the Mobile Network Middleware and the Virtual Organisation Management. Where, by means of a user identity model and user identity management framework a dynamic establishment of administrative domains commercially operated by one operator to a VO based on the user-centric needs is addressed. Virtual Organisations having as motivation to provide the means for ubiquitous collaboration among mobile partners will face the need to select the necessary components from every layer in order to perform it efficiently. This is subject not only to the adopted policy for VO management and the selected business models, but also to constraints and/or facts originating from the networking level: wireless hot spots can make use of peer-to-peer computing and ignore access to the legacy networks. One new aspect, which is currently not considered sufficiently in the existing approaches in the grid community is the new requirement, that mobile users as part of the VO might disappear temporarily. For example, a user sitting in a train passes a tunnel and has for 30 seconds no network connection and when leaving the tunnel the same user is connected via a new operator to the same VO. This requires again a coherent management in order to allow the user to continue the session where it was stopped without significant overhead. The FeGIMa architecture will support connectivity seamlessly, which is currently not supported in existing approaches.

The orchestration of services in the mobile Grid environment implies a vertical approach, since the main aspect that has to be addressed is the adaptation and Configuration of services, resources, access rights (security) and networking attributes throughout the various layers of the FeGIMa platform. Orchestration makes essential the use of application-specific functionalities and the adoption

of advanced brokering schemes to be used for the final synthesis of services. This synthesis is motivated by a session-based approach which will make obvious that the provision of the Application Specific Service to the end user will be subject to a broader configuration management scheme of the underlying modules and sub-modules. Orchestration and Configuration Management is clearly an essential element of FeGIMa as the enabler of dynamic applications. The usage of orchestrated dynamic applications forms central importance to FeGIMa as it allows adapting the application on changing situations.

The orchestration of Web Services is already moving fast and several competing specifications have emerged. Most noticeable the BPEL4WS specification under the auspices of OASIS, the W3C Choreography Working Group and also the new WS-CAF Framework presented by an industrial consortium. However as Web Services are stateless and the way data is communicated between dependant services on a workflow differs substantially from the needs for Grid Services noted in vertical issues of FeGIMa above. Another issue is the transient nature of Web Services that require the support of on-the-fly instantiation in workflow description languages. Beside the need for notification towards the workflow engine that due to a change of the administrative domain a new set-up is required that is not addressed at all in existing specifications or products. Also during the orchestration process, a close interaction between the network middleware layer down to the network layer is required in order to realise movements of a mobile end-system requesting a service/orchestration process and to adapt in near-real-time to the changing network conditions appropriately

5. Conclusion

The results of the FeGIMa project will enable the next generation of grid middleware one should no longer rely on the basic functionality of the traditional Internet. It should foster the utilisation of the new capabilities of the next generation Internet designed with the mobile and roaming user in mind. Simply using the mobile Internet is not enough because of various limitations partially described in the previous section. One has to establish communication points between the different layers identified and enhance their functionality in order to serve the upper layers in a more efficient way.

Current Internet philosophy does not consider commercial facets. In fact there exists no operational

concept to address how roaming agreements and the business relationship between the service providers (network, roaming, Grid Services, Information provided through Grid and Web Services, etc.) and the service consumer with several parties involved in the service provisioning process. Assuming that telecom operators limit themselves through specific service locators the identification process of potential services contract or payment based service access can be established. The service locator could use the user profile provided by his home AAA server in order to identify the services that can be offered to the user in accordance to the creditability and contract the user has with his primary operator. This could be one approach for commercialisation of the value chain from client over the network operator up to the service provider that will be investigated by the consortium.

The explosion of Grid projects and applications world-wide has led to a diversity of approaches. Some Grid computing toolkits are widely used, but none of them has universal acceptance. All of them rely on the standard Internet widely deployed and available almost everywhere. The Internet as it is so far offers only basic transport services and can't offer sophisticated support for fundamental properties of an application middle-ware such as user identification and authorisation or the support for commercial exploitation of offered services.

6. References

- [1] Liberty Project www.projectliberty.org
- [2] Internet 2; <http://www.internet2.edu/>
- [3] European Data Grid (EDG) <http://eu-datagrid.web.cern.ch/eu-datagrid/>
- [4] Shibboleth <http://shibboleth.internet2.edu/shib-intro.html>
- [5] Security Assertion Markup Language (SAML) <http://www.oasis-open.org/specs/index.php#samlv1.1>
- [6] Kerberos <http://gost.isi.edu/info/kerberos/>
- [7] Private Key Infrastructure (PKI) <http://csrc.nist.gov/pki/>
- [8] Transport Layer Security <http://www.consensus.com/ietf-tls/ietf-tls-home.html>