# D6.1 Report on the Assessment of Operational Procedures and Definition of the MAPPER Operational Model

Project acronym: *MAPPER*

Project full title: Multiscale Applications on European e-Infrastructures.

Grant agreement no.: 261507

| | |
|---|---|
| **Due-Date:** | Month 6 |
| **Delivery:** | Month 6 |
| **Lead Partner:** | LMU |
| **Dissemination Level:** | PU |
| **Status:** | Final |
| **Approved:** | Q Board, Project Steering Group |
| **Version:** | 1.1 |

## DOCUMENT INFO

| Date and version number | Author | Comments |
|---|---|---|
| 17.03.2011 v1.0 | Ilya Saverchenko | Final draft |
| 31.03.2011 v1.1 | Ilya Saverchenko | Updated the document based on comments provided by Krzysztof Kurowski and Stefan Zasada |

**TABLE OF CONTENTS**

# 1 Executive summary

This deliverable is a living document reporting on the assessment of operational procedures of the European e-Infrastructures targeted by MAPPER, including DEISA, EGI and PRACE, and defining the MAPPER operational model. This deliverable describes the status and intermediate results achieved by the tasks 6.1, Coordination of Operation and Support, and 6.3, VO Management and Support. Subsequent versions of this deliverable are expected to contain a summary of the outcomes of task 6.3 and reference the deliverable D6.3, Support Process Definition, for the detailed description of the task's status. During the first 6 months WP6 was primarily focused on the assessment of operational procedures. As such the first issue of the deliverables D6.1 and D6.3 are combined in this document. It is expected that following issues of the deliverables will be separated with the first release of deliverable D6.3 available in month 12.

The first issue of deliverable D6.1 focuses on assessment of operational procedures implemented by the targeted European e-Infrastructures, including DEISA, EGI and PRACE. The assessment process was carried out in relation to the objectives of the MAPPER project and focused on operational procedures that affect service provisioning and evolution of the e-Infrastructures from the perspective of the scientific community. In the first phase of the project policies were identified and selected for the assessment affecting the following three areas: Authentication, Authorisation and Accounting, Network and Compute. It is expected that this list will be re-evaluated and extended in the subsequent phases of the project in accordance with the requirements of the MAPPER research infrastructure.

# 2 Introduction

One of the main goals of the MAPPER project is the implementation and operation of a distributed research infrastructure capable of supporting execution of complex multiscale simulations. To be successful in this challenging task MAPPER partners have to define and implement the organisational processes, policies and agreements necessary for coordinated operation of the infrastructure.

The focus of WP6 is the definition of the MAPPER operational model, deployment and operation of the MAPPER research infrastructure and support of the research communities. It is expected that the MAPPER operational model will facilitate integration with production European e-infrastructures, such as DEISA, EGI and PRACE. Operational policies adopted by these projects will be thoroughly studied and used as a reference for refinement and adaptation of the MAPPER operational model.

This document consists of two parts: assessment of operational procedures defined in the European e-Infrastructures and definition of the MAPPER operational model. The document will be regularly updated to reflect the current state of the MAPPER and European research infrastructures.

# 3 Assessment of Operational Procedures

## 3.1 Overview of European e-Infrastructures

The European distributed computing landscape is shaped by the three biggest e-Infrastructure: DEISA [1], EGI [2] and PRACE [3], that will be introduced in the following sections. Each e-Infrastructure addresses the needs of specific research groups. PRACE forms the top layer of the European distributed computing ecosystem and operates state of the art peta- and exascale High Performance Computing resources. DEISA acts as the middle layer and federates national European HPC systems. EGI, established the bottom layer of the ecosystem, brings together national and regional providers to enable collaboration of research communities across Europe. Together the projects complement each other and provide a complete set of computing services to the European scientific community.

MAPPER will focus on integrating its services in the European e-Infrastructure listed above. Therefore it is important for MAPPER to study, analyse and adopt operational policies and procedures defined by DEISA, EGI and PRACE.

### 3.1.1 PRACE

The Partnership for Advanced Computing in Europe, PRACE, is a unique persistent pan-European Research Infrastructure for High Performance Computing. PRACE provides Europe with world-class systems for world-class science and strengthens Europe's scientific and industrial competitiveness. PRACE maintains a pan-European HPC service consisting of up to six top of the line leadership systems well integrated into the European HPC ecosystem. Each system provides computing power of several Petaflop/s (one quadrillion operations per second) in midterm. On the longer term Exaflop/s (one quintillion) computing power will be targeted by PRACE. This infrastructure is managed as a single European entity.

The partnership was established through the close collaboration of the European countries that prepared the legal, financial, and technical basis of the project. The First Implementation Phase of PRACE is in line with the objectives of the PRACE Research Infrastructure organisation: from coordinated system selection and design, coherent management of the distributed infrastructure, software deployment, porting, scaling, optimising applications and promoting and advancing application development and the skills.

### 3.1.2 DEISA

DEISA, the Distributed European Infrastructure for Supercomputing Applications, is a consortium of leading national supercomputing centres that aims to foster pan-European world-leading computational science research.

DEISA deploys and operates a persistent, production quality, distributed supercomputing environment with continental scope. It aims at delivering operational solutions for a future European HPC eco-system. By extending the European collaborative environment in the area of supercomputing, DEISA is paving the way towards the deployment and operation of a persistent cooperative European HPC ecosystem, as suggested by ESFRI [4].

### 3.1.3 EGI

Building a world-class pan-European High Performance Computing Service and infrastructure involves the scientific and industrial user communities with their leading edge applications. This needs to be done in a rapidly evolving context, where technologies change continuously and where the science focus changes as results are obtained and new directions are explored.

The ultimate goal of EGI is to provide European scientists and their international partners with a sustainable, reliable e-Infrastructure that can support their needs for large-scale data analysis. This is essential in order to solve the big questions facing science today, and in the decades to come.

EGI will collect user requirements and provide support for current and potential new user communities, for example the ESFRI projects. The project will also support the current heavy users of the infrastructure, such as high energy physics, computational chemistry and life sciences, as they move their critical services and tools from a centralised support model to one driven by their own individual communities.

### 3.1.4 PL-GRID

PL-Grid operates the National Grid Initiative (NGI) in Poland. PL-Grid aims at significantly extending the amount of computing resources provided to the Polish scientific community and constructing a Grid system that facilitates effective and innovative use of the available resources. In this aspect, we focus on exploitation of computational and storage facilities by virtual organizations and on implementing a comprehensive Grid resource management suite, comprising approximately 35 individual services. PL-Grid draws upon the experience of European initiatives, such as DEISA, EGI and PRACE, the scientific results attained by individual partners and the outcome of R&D activities carried out within the project.

In the framework of cooperation with EGI, PL-Grid project is involved in:

- EGI.eu council activities through membership in the EGI.eu Executive Board EGI-InSPIRE.
- Global and International tasks including operation and oversight of the EGI e-Infrastructure and coordination of resource allocation and brokering support for NGI VOs'.
- Organization and management of the computational chemistry and material science communities.
- Development of unified middleware via the European Middleware Initiative.
- Scientific application porting within the EGI Application Porting Specialised Support Centre.
- Providing EGI service interfaces for the availability monitoring, issue tracking, user support and accounting systems.

## 3.2 Assessment of Authentication, Authorisation and Accounting Policies

Authentication, authorization and accounting refer to processes employed by resource providers, such as e-Infrastructures, to regulate access and usage of their resources and services.

Authentication describes the process of validating user's digital credentials for establishing identity of the user. One of the most well known authentication methods is based on a user name and a password, where the user name identifies the user and the password is used to prove the identity. Although it is one the most common authentication methods, it has a number of drawbacks and therefore is not as interesting for MAPPER as, for example, the X.509 certificate-based authentication. The latter authentication method relies on the Public Key Infrastructure scheme that is well established across Europe and offers many advanced features that greatly simplify the authentication process.

The aim of authorization is to determine whether the user is allowed to perform specific functions, such as the access and use of resources and services. For example, many European e-Infrastructures allow users to access only a subset of available computing resources. Therefore when a user tries to access a resource, the user's identity is checked to determine whether access can be granted, i.e. whether the user is authorized to access this resource.

Accounting is the process of collecting and analysing data that describes resource consumption by users. For instance, a resource provider might collect information about service usage in order to prevent monopolization of the service by a single community or user. The e-Infrastructures targeted by MAPPER often try to balance resource usage by different research groups by allocating fixed budgets for storage, computing time and other resources.

### 3.2.1 PRACE

To get access to the PRACE e-Infrastructure, research groups have to submit project proposals describing resource requirements, methods and models used, research goals, and the scientific merit of their work. The proposals are evaluated by a committee composed of PRACE representatives and external experts based on a set of predefined criteria. After the top proposals are selected a project is created and assigned to each of the respective research groups. PRACE takes over management of the projects including such activities as the allocation of resources, creation of user accounts and maintenance of authorization and accounting facilities. This allows the users to focus on the research but, at the same time, introduces some constrains since user requests can be processed only by PRACE representatives.

To access PRACE resources and services each user has to authenticate him- or herself. The authentication process adopted by PRACE is based on the PKI scheme. As mentioned above the scheme is well established across Europe. The majority of research facilities in France, Germany, or the Netherlands, for instance, are capable of issuing X.509 certificates that can be used to authenticate to PRACE resources. However, to get a certificate a person has to be physically present at the so-called Registration Authority. In cases when this is inconvenient or not possible, a number of alternative schemes are available.

The certification process is not as seamless across all Europe and in some cases might require more time and effort. To circumvent this group certificates shared by several users could be used. Unfortunately, this is not possible since the PRACE policy demands that a certificate uniquely identifies a user.

As a part of the project definition process, each research group is allocated a fixed budget for computing, storage and other services and is granted access to one PRACE resource. Each project has a limited lifetime that normally does not exceed one year. This introduces some limitations since within the given time frame each group has to focus on achieving their scientific objectives. It is expected that the major part of all work is carried out on the assigned PRACE resource. As such, access to resources or services hosted by other PRACE partners, with the exception of central services, is not possible under normal circumstances.

Each PRACE user has access to the accounting information, e.g. computing time and disk quota, for his or her account. One person in each research group assumes the role of Principal Investigator, which grants that person access to the accounting information for the whole group.

### 3.2.2 DEISA

DEISA AAA policies are similar to the PRACE policies described in the previous section, but there are a few exceptions. DEISA offers more flexibility with regard to access to the resources and services available in the e-Infrastructure. Each research group can get access to any DEISA site and therefore all resources and services hosted by the site upon request. Research groups are not allowed to exceed their allocated budget but in many cases it is possible to relocate a part of the remaining budget from one DEISA resource to another.

In addition to the concept of projects, DEISA supports so-called virtual communities. A virtual community is an entity that federates several research groups working in the same scientific area, such as biosciences or plasma and particle physics. Upon creation, a virtual community is allocated an initial set of resources. The resources then have to be split among one or several projects defined by the members of the community. As such virtual communities simplify the resource allocation process. However all virtual community projects have to follow the DEISA AAA policies that apply to the regular projects.

Virtual communities offer users additional flexibility by delegating some of the management tasks to its members by making it easier to define new and manage existing community projects, relocate parts of the community budget from one project to another, extend the lifetime of a project and so on. Yet, unfortunately, virtual communities are not able to remove the majority of limitations introduced by the DEISA AAA policies.

### 3.2.3 EGI

Access to EGI resources is granted to users based on their virtual organization membership, where a VO is a dynamic set of individuals and institutions active in a specific scientific area. From this point of view VOs are similar to virtual communities defined in DEISA. Yet a VO, generally speaking, has more responsibility and control over users, resources and services belonging to the VO. As such EGI defines generic policies for authentication, authorization and accounting. Each VO is able to define and follows its internal rules as long as they are inline with EGI.

Similarly to DEISA and PRACE, EGI relies on the PKI scheme for authentication. This means that users require an X.509 certificate to gain access to EGI resources. Yet, unlike the other European e-Infrastructures, EGI does not enforce the one-to-one mapping for users, user accounts and certificates. This means that several users are allowed to share a common certificate associated to a pool account and use it for authentication with EGI resources and services. This is a very useful and desired feature since it allows VOs to setup internal authentication mechanisms for simplifying the authentication process, for example in cases when not everyone who needs to access the e-Infrastructure is able to get a personal certificate. Regular user accounts requiring a personalized certificate are, of course, supported as well.

EGI does not define authorization and accounting policies and delegates this task to the individual VOs. As such, the VOs are able to internally control access to resources and services and specify mechanisms used for accounting. Policies defined by each VO should be inline with the local legislative regulations. For instance, resource providers should handle user related data according to the locally enforced laws.

### 3.2.4 Pl-GRID

Similar to EGI and PRACE, PL-Grid relies on the PKI scheme for authentication. This means that users require an X.509 certificate to gain access to PL-Grid resources.

## 3.3 Assessment of Network Policies

A high quality network connection is essential for multiscale simulations spanning across multiple geographical locations. Poor network performance caused by, for example, high latency or random jitter, can compromise the ability of the individual simulation components to exchange data which may result in a failure of the simulation. As such it is important that network policies enforced in an e-Infrastructure offer users the possibility to satisfy their network requirements either by using available resources and services or by requesting support from external providers.

### 3.3.1 PRACE

PRACE resources are interconnected with a dedicated high-speed 10Gb network that is shared by all users. Although offering a very good performance in theory, it can not always be achieved in practice, such as in situations where many users are concurrently transferring large amount of data. At the moment PRACE operates only two machines, JUGENE located in Forschungszentrum Jülich and CURIE operated by the French Alternative Energies and Atomic Energy Commission. As such, utilization of the PRACE internal network is not high at present. This will however change later this year when additional systems will be brought in production.

Connection to the public network is available on all PRACE resources and is managed by the resource provider. PRACE does not define policies for this type of connection and therefore their availability might vary from one resource to another.

### 3.3.2 DEISA

The majority of DEISA resources are interconnected with a dedicated high-speed 10Gb network. In a few cases where a high-speed interface cannot be installed due to technical limitations a 1Gb link is provided. Availability and performance of all network interfaces and segments is monitored by DEISA and the respective network operators to ensure high quality of service provided to the users. The DEISA internal network is shared and can be concurrently used by all DEISA users. Therefore all limitations described in the previous section apply to DEISA.

Due to the heterogeneous nature of the DEISA e-Infrastructure the network performance optimization capabilities are technically limited. As a consequence, in some cases research groups might not be able to utilize the network to its full potential. To circumvent this a research group may request a dedicated connection between specific DEISA resources. Such network links can be provided by the National Research and Education Networks operating in the regions respective to the resource location. Depending on the agreement between DEISA and the NRENs involved the research group issuing the request might be asked to bear a part of the outstanding costs.

### 3.3.3 EGI

NGIs and resource providers are responsible for establishing the necessary network connectivity using the public Internet. Each NGI and VO may establish internal regulations and policies that, however, will not be uniform across the EGI e-Infrastructure. Research groups that require special networking functionality, such as links with low latency for computational steering workflows, may request dedicated links from the respective NREN or the DANTE project. However, as in the example describe in the previous section, the requester might be asked to bear a part of the costs associated with setup and maintenance of the network channel.

### 3.3.4 PL-GRID

The PL-Grid resource providers are connected together with the PIONIER National Research Network. The PIONIER network is a nationwide broadband optical network providing a base for research and development in the areas of information technology, telecommunications and computing sciences. Built entirely from the Committee for Scientific Research funds, the network connects 21 Academic Network and 5 HPC Centers. An important element of the PIONIER network is the cross border connection to other European network operators. This allows establishing cheap and quick access to major European Internet Exchange Points (IXPs) and connect to European National Research and Education Networks (NRENs).

## 3.4 Assessment of Compute Policies

Compute policies define how e-Infrastructure users interact with available computing resources. This includes, for example, submission and management of batch jobs, simulation steering and workflow management. Functionalities such as co-allocation and advance reservation greatly depend on the policies as well. Therefore it is essential that the compute policies enforced on the European e-Infrastructures are evaluated during the course of MAPPER as they have a direct impact on the progress and outcome of the project.

### 3.4.1 PRACE

PRACE offers several different ways to access its compute resources including interactive access, local job scheduler systems and a unified interface for job submission. This allows PRACE to offer several alternative ways for management of compute jobs, which is beneficial for the user. Yet, at this moment, advanced functionality such as advance reservation of compute resources is not available. As a results research groups working on the PRACE e-Infrastructure are not able to efficiently plan distribution of work, for instance submission of simulations with multiple components, which often results in delays and the inability to spend the allocated budget during the tight time frame set by the project lifetime.

Due to the HPC nature of the e-Infrastructure, PRACE user projects are assigned a single compute resource. It is therefore expected that users work only on the allocated machine and do not execute simulations that interact with external services or components.

Another limitation comes from the fact that PRACE users are not allocated special batch queues or pools on PRACE compute resources and have to share the machines with other users. Because of this it is normally hard to estimate when the necessary resources will become available. This does not affect single simulation runs that are independent of one another. However it prevents users from running multiscale simulations that cross e-Infrastructure boundaries resulting in limited interoperability of PRACE with other European e-Infrastructures.

### 3.4.2 DEISA

DEISA compute policies are, to a large extent, similar to the ones defined in PRACE. As well as PRACE, DEISA offers a number of different ways for users to access its compute resources. And, although no advanced resource management functionality is available, the larger total number of resources and technologies provides users with more alternatives. For instance, DEISA partners are able to provide research groups with dedicated or prioritized access to resources, which, under certain conditions, can be used to replace such functionality as advance reservation.

### 3.4.3 EGI

EGI does not define any compute policies and delegates this task to the individual VOs. This strategy does not harm the interests of the scientific communities, as each VO representing a research area is able to agree on a suitable set of policies that apply to all resources available in the VO. This however might introduce unexpected complication for the deployment of MAPPER services since the deployment strategy and the respective policies will have to be discussed separately with each VO.

### 3.4.4 PL-GRID

PL-Grid implements grid operation according to a novel resource allocation centric architecture. Compute resources are allocated according an SLA signed between users and respective resource providers.
The main actors involved in the resource allocation process are:

- VOs and users groups use assigned resources in computational experiments.
- Resource Providers (RP) need to plan efficient usage of their resources.
- National Grid Initiative brokers resources provided by RPs. NGI can act as proxy provider offering resources to the users, and as proxy user for some RPs; a similar role can be applied to others NGIs in European Grid Initiative.

# 4 MAPPER Operational Model

The MAPPER operational model should define organisational structures, processes and agreements among partners necessary for operation and support of the MAPPER research infrastructures. The operational model is divided in to three areas: deployment of the MAPPER infrastructure components, monitoring and quality assurance and VO and user support, all of which are introduced and described below.

## 4.1 Deployment of the MAPPER Infrastructure Components

Prior to deployment in the MAPPER e-Infrastructure, all software components have to pass a rigorous validation and testing process performed by WP4 and WP5. The tests have to confirm that the software to be deployed fulfils the functional requirements defined by the developers and is able to interact with the environment, for instance the operating system, drivers and low level libraries and applications, without disturbing production or affecting systems performance in any undesired way. During the later stages of the project the MAPPER components will be thoroughly tested to ensure their interoperability with each other.

To support the software validation process a test bed environment has been created. The environment is meant to act as a platform for testing activities and experimentation with the MAPPER software components, for instance the scientific applications developed by WP7 and WP8. Production quality resources for the test bed were provided by D-Grid, PL-Grid and NGS, the national Grid projects of Germany, Poland and the United Kingdom respectively.

All successfully validated software components should be passed to WP6 for deployment on the MAPPER research infrastructure. Each component will be assigned a WP6 service expert responsible for co-ordinating its installation, configuration and maintenance. To facilitate the deployment process WP6 developed the MAPPER Service Description Form. The form is used to describe the software components and provide their essential characteristics including system and security requirements. It is expected that the form will be filled out for each of the components to be deployed and provided to the responsible application expert who will use the information during the initial installation phase.

After a component is successfully deployed on the MAPPER resources the service expert will verify its proper functionality and ensure that maintenance activities, for instance installation of necessary updates, are performed regularly.

## 4.2 Monitoring and Quality Assurance

All deployed MAPPER infrastructure components will be monitored to ensure their availability and correct functionality. Monitoring will be performed using applications and tools selected by the MAPPER Technical Board according to functional requirements devised by the Technical and Application Boards. During the selection process the MAPPER Technical Board will pay especial attention to the monitoring solutions used by DEISA, EGI and PRACE to guarantee full compatibility with European e-Infrastructures in this area. Recommendations of the Technical Board will be regularly presented to WP6, which will carry out deployment and maintenance of the proposed applications and tools on the MAPPER e-Infrastructure.

The primary goal of the monitoring efforts is to collect information about availability of the MAPPER infrastructure components. Therefore initially WP6 will focus on service-level monitoring and deploy tools that can be used to check whether MAPPER services are operational and respond to basic queries and requests. Such tests will include, for instance, checks for service processes, remote connection to service interfaces and analysis of log files. It is expected that members of MAPPER service and joint research work packages will provide feedback and suggestions on the state of the monitoring activities. This input will be used by WP6 to extend functionality and improve service quality by deploying additional monitoring solutions, for example tools for user-level monitoring, and developing new monitoring probes.

Detected incidents and unexpected behaviour of the MAPPER infrastructure components will be directly reported to WP6. Each incident will be classified based on its impact on the MAPPER e-Infrastructure and assigned to the service expert responsible for the failing component. The expert will analyse and attempt to resolve the problem. If successful, all parties affected by the problem will be informed and the solution will be documented in the MAPPER knowledge base. In situations when members of WP6 cannot resolve a problem, WP4 and WP5 will be asked for assistance. If necessary, the problem will be escalated to the developers of the respective component.

All monitoring data collected by MAPPER will be archived and used for reporting and quality assurance. General information about availability and functionality of MAPPER services will be published on the MAPPER web site and used for dissemination and reporting. The production status of the MAPPER infrastructure components will be regularly evaluated by the project partners, in particular by the MAPPER Technical Board. The board will closely observe the service quality levels and the overall availability of the MAPPER e-infrastructure and will issue recommendations to WP6 if the values fall below defined levels.

## 4.3 VO and User Support

Over the course of the project MAPPER will offer its services and support multiple VOs. The support activities will start with operation of the MAPPER fast track components and will be provided in a number of ways, including documentation, training, VO management and other services, to assure high quality of service and sustainability of the MAPPER research infrastructure.

VO management covers the whole life cycle of a VO starting from its creation through operation until a concerted termination. It includes such tasks as authentication, authorization, certificate management, accounting and billing. MAPPER will support two classes of VOs: one for users interested in testing and experimentation and another for communities wishing to use MAPPER services to support their research activities. WP6 will provide full support to both classes and will define policies for user migration from testing to production VOs.

Comprehensive documentation, including but not limited to installation and administration manuals, user guides, tutorials and sample videos, will be offered for all MAPPER infrastructure component. These materials will focus on achieving the following goals. First, providing detailed information and guidance to MAPPER users who want to run complex multi-scale simulations. And second, acting as a reference for e-Infrastructure operators charged with installation and configuration of MAPPER components on production e-Infrastructures. The documentation will be regularly improved and extended and will play an important role in achieving the sustainability goals of the MAPPER project.

Introductory lectures and training course will be provided as a part of the MAPPER dissemination activities led by WP2 and supported by all project partners. The main focus of these activities will be to inform user communities about MAPPER services and help them in the initial porting and enabling of scientific workflows on the MAPPER research infrastructure.

MAPPER will operate a help desk and provide VOs and user communities with comprehensive technical support. Reported problems will be handled by WP6 in close co-operation with all other work packages according to the ISO 20000 standard. Issues beyond the expertise of MAPPER will be escalated to external entities, such as application developers or service providers.

# 5 Conclusion

This document outlines the MAPPER operational model and describes its key characteristics including deployment and operation of MAPPER fast and deep track components, monitoring and quality assurance of the provided services and user support. The model will be regularly re-assessed and adopted to satisfy the evolving requirements of the MAPPER infrastructure. Changes and development of the MAPPER operational model will be reflected in the future issues of this deliverable.

# 6 Acronyms and References

## 6.1 Acronyms

| | |
|---|---|
| AAA | Authentication, Autorization, Accounting |
| DANTE | Delivery of Advanced Network Technology to Europe |
| DEISA | Distributed European Infrastructure for Supercomputing Applications |
| EGI | European Grid Initiative |

| | |
|---|---|
| ESFRI | European Strategy Forum on Research Infrastructures |
| HPC | High Performance Computing |
| ISO | International Organization for Standardization |
| NGI | National Grid Initiatives |
| NREN | National Research and Education Network |
| PKI | Public key infrastructure |
| PRACE | Partnership for Advanced Computing in Europe |
| VO | Virtual Organisation |
| WP | Work Package |

## 6.2 References

[1] DEISA http://www.deisa.eu/

[2] EGI http://www.egi.eu/

[3] PRACE http://www.prace-project.eu/

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