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 24
Delivery:	Month 24
Lead Partner:	LMU
Dissemination Level:	PU
Status:	Final
Approved:	Q Board, Project Steering Group
Version:	2.1

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
15.09.2012 v2.0	Ilya Saverchenko	Updated version of the document
27.09.2012 v2.1	Krzysztof Kurowski	Minor changes and updates to the document

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 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 Summary of changes

The following changes were made to this document since the last version 1.1:

1. Updated description of European e-Infrastructures to reflect transition from DEISA to PRACE.
2. Expanded and improved assessment of AAA, network and compute policies.
3. Added an assessment of user support policies.
4. Updated MAPPER operational model

3 Introduction

One of the main goals of the MAPPER project is the implementation and operation of a distributed computing 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 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 initial definition of the MAPPER operational model. The document will be regularly updated to reflect the current state of the MAPPER and European e-infrastructures.

4 Assessment of operational procedures

4.1 Overview of European e-Infrastructures

The European distributed computing landscape is shaped by two biggest e-Infrastructure: EGI [2] and PRACE [3], that will be introduced in the following sections. Each e-Infrastructure addresses the needs of specific research groups, in particular specific resource requirements of applications. PRACE comprises the top two layers of the European distributed computing ecosystem and distinguishes two classes of High Performance Computing resources: Tier-0 and Tier-1. PRACE Tier-0 operates state of the art peta- and exascale (by 2020) European HPC resources. PRACE Tier-1 acts as the middle layer and federates national European HPC systems. Operational policies adopted by PRACE Tier-0 and Tier-1 infrastructures differ slightly. PRACE Tier-1 is a follow up to the DEISA [1] project, thus inheriting operational

policies and procedures. PRACE Tier-0 on the other hand adopts a set of unique policies to reflect requirements of European peta- and exascale HPC systems in the future.
EGI establishes the bottom layer of the ecosystem and 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 various European scientific communities.

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 EGI and PRACE. Additionally, it is assumed that the MAPPER operational model will influence and hopefully improve the way computing services will be delivered and supported by EGI and PRACE, especially for research communities interested in distributed multi-scale simulations.

4.1.1 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. Federating HPC systems on European and national levels allows PRACE to provide state of the art computing services to European scientific communities.

PRACE maintains a pan-European HPC service consisting today of six top of the line Tier-0 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 by 2020. This e-infrastructure is managed as a single European entity.

PRACE Tier-1 integrates national HPC resources to provide all members of the European scientific community with access to high-class HPC systems. PRACE deploys and operates a persistent, production quality, distributed computing infrastructure 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, PRACE is paving the way towards the deployment and operation of a persistent cooperative European HPC ecosystem, as suggested by ESFRI [4].

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 and Second Implementation Phases of PRACE are in line with the objectives of the PRACE Research Infrastructure organisation: from coordinated system selection and design, coherent management of the distributed computing infrastructure, software deployment, porting, scaling, optimising applications and promoting and advancing application development and end-users skills.

4.1.2 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.

4.1.3 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 distributed computing resources located in five supercomputign centers in Poland. 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 PL-Grid project and its follow up called PL-Grid Plus enganing directly fourteen different scientific communities.

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.

4.2 Assessment of authentication, authorisation and accounting policies

Authentication, authorization and accounting refer to processes employed by resource providers, for instance e-Infrastructures, to regulate access, usage and report on cost of their resources and services.

Authentication describes the process of validating user's digital credentials for establishing identity of the end-user. One of the most well known authentication methods is based on a user name and a password, where the user name identifies the end-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 at this stage 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 an end-user is allowed to perform specific functions, such as access and use of resources and services. For example, many European

e-Infrastructures allow end-users to access only a subset of available computing resources. Therefore, when the end-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 end-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 the end-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.

4.2.1 PRACE Tier-0

To get access to the PRACE Tier-0 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 appointed by the PRACE AISBL[]. The PRACE Scientific Steering committee includes invited representatives and external experts representing diverse scientific domains. The goal of the committee is to select top proposals capable of leveraging capabilities of available HPC resources to their full extent.

Each accepted proposal is allocated a project assigned to one of the production Tier-0 systems. PRACE is responsible for management of the projects including allocation of resources, creation of user accounts, management of access credentials and maintenance of authorization and accounting facilities. This allows end-users to focus on the research but, at the same time, introduces constrains since user requests can be processed only by PRACE representatives.

To access PRACE resources and services each end-user has to authenticate him- or herself. The authentication process adopted by PRACE Tier-0 is a combination of traditional user name and password based access and more modern X.509 certificate PKI scheme. Combination of both methods is necessary to accommodate varying functional and security requirements of Tier-0 resources.

The user name and password scheme is well established and is generally available on production systems. As such, from the operations perspective, it offers itself as the

straightforward choice for end-user authentication. This method, however, is coupled with a number of limitations. It is not always possible to ensure security and privacy during the initial distribution of credentials. Moreover it offers limited control and auditing what potentially hinders security.

The alternative adopted by many PRACE partners is the X.509 certificate PKI authentication scheme. The scheme relies on secure PKI credentials for authentication. The credentials comprise public and private keys, where, as can be seen from the designation, the public key is open while the private is kept secure by the end-user. The scheme is well established across Europe. The majority of research facilities in France, Germany, the Netherlands, Spain, Switzerland and other countries operate infrastructures for management of X.509 certificates. Certificates issued by these authorities follow international standard defined by organisations such as EUGridPMA and are recognised and accepted by all PRACE partners.

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 Tier-0 resource. Each project has a limited lifetime that normally does not exceed one year. This might lead to restrictions since scientific goals and objectives set for a project should be achieved within the given time frame. 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.

All end-users have access to resource accounting information, e.g. computing time and disk quota, for his or her account. One or several persons in each research group assume a role of Principal Investigator and are granted access to the accounting information for the whole project.

4.2.2 PRACE Tier-1

PRACE Tier-1 policies are similar to those adopted by Tier-0 systems with a few notable exceptions. PRACE Tier-1 offers more flexibility with regard to access to the resources and services available in the e-Infrastructure. Each research group is granted access to multiple resources. Furthermore access to additional systems can be granted during a lifetime of a project 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 PRACE Tier-1 resource to another.

In addition to the concept of projects, PRACE Tier-1 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 PRACE 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 PRACE AAA policies.

4.2.3 EGI

Access to EGI resources is granted to users based on their virtual organization membership, where a VO is a dynamic entity that brings together individuals and institutions active in a specific scientific area. From this point of view VOs are similar to virtual communities defined in PRACE. 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 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 to access EGI resources and services. This is a 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.

4.2.4 PL-Grid and PL-Grid Plus

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. However, the overall process of requesting and managing X.509 certificates has been simplified by offering web-based tools to end-users.

Resource, service access and accounting policies enforced on PL-Grid resources are similar to those adopted by EGI.

4.3 Assessment of network policies

A high quality network connection is essential for distributed multiscale simulations spanning across multiple geographical locations. Poor network performance caused by, for example, high latency or random jitter, can compromise the ability of individual models in a multiscale simulation to exchange data what may result in a failure of the simulation itself. As such it is important that network policies enforced within and between e-Infrastructures meet potential end-user requirements regarding network connection properties and performance.

4.3.1 PRACE Tier-0

PRACE Tier-0 resources are interconnected with a dedicated high-speed 10Gb network that is shared by all end users. Although offering a very good performance in theory, it cannot always be achieved in practice as real-time performance depends on a range of factors that cannot be foreseen in advance. At the moment PRACE Tier-0 operates four machines, CURIE operated by the French Alternative Energies and Atomic Energy Commission, Hermit hosted by the High Performance Computing Center Stuttgart, JUGENE located in Forschungszentrum Jülich and SuperMUC of Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities. At the moment utilization of the dedicated network at the Tier-0 level is limited. End-user projects are granted access to single systems only, thus the dedicated network is primarily used for data transfers to and outside of the PRACE Tier-0 ecosystem. This situation is expected to change after PRACE Tier-0 access model evolves.

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.

4.3.2 PRACE Tier-1

The majority of PRACE Tier-1 resources are interconnected with a dedicated high-speed 10Gb network. In a few cases where a high-speed interface cannot be installed for technical or financial reasons a 1Gb link or virtual interface over public network is provided. Availability and performance of all network interfaces and segments is monitored by PRACE or network providers, primarily the GÉANT project, to ensure high availability and quality of services offered to end-users. The PRACE internal network is shared and can be concurrently used by all PRACE end-users. Therefore limitations described in the previous section apply to PRACE Tier-1 as well.

Due to the heterogeneous nature of the PRACE 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 PRACE resources. Reserved network links can be provided by National Research and Education Networks operating in the regions where respective resources are located. Depending on the agreement between PRACE and the NRENs involved the research group issuing a request might be asked to bear a part of the outstanding costs.

4.3.3 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 and GÉANT projects. 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.

4.3.4 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 Networks and 5 HPC Centres. 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 NRENs.

4.4 Assessment of compute policies

Compute policies define how e-Infrastructure users interact with available computing resources. This includes interactive access to computing resources, submission and management of batch jobs, advance reservation and co-allocation of computing resources, simulation steering and workflow management. It is essential that 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.

4.4.1 PRACE Tier-0

PRACE Tier-0 offers several different ways to access its compute resources including interactive access, local job scheduler systems and a unified interface for job submission. These services are described in PRACE service catalogue and are available on all PRACE Tier-0 resources. 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 distribute and schedule their work. For instance, submission of simulations with multiple components often results in delays and the inability to spend the allocated budget during the tight time frame set by the project lifetime.

Due to source allocation policies, PRACE end-user projects are assigned a single HPC resource. It is therefore expected that users work only on the allocated machine and do not need to execute simulations that interact with external services or components.

Another limitation comes from the fact that PRACE end-users are not allocated special batch queues or pools on PRACE compute resources and have to share machines with local 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.

4.4.2 PRACE Tier-1

PRACE Tier-1 compute policies are, to a large extent, similar to the ones defined for PRACE Tier-0 systems. RACE Tier-1 offers a number of different ways for users to access its compute resources extending supported access protocols and endpoints beyond those defined on the Tier-0 level. Furthermore each PRACE Tier-1 partners is entitled to provide and support custom resource access methods.

Advance resource management services are not a part of the PRACE service catalogue. Through tight collaboration between MAPPER and individual PRACE Tier-1 partners we were able to initiate efforts in setting up and testing advance reservation. As of summer 2012 two PRACE Tier-1 partners: EPCC and SARA, support advance reservation of computing resources in a semi-automated manner. Advance reservations can be made by end-users on EPCC HECToR and SARA Huygens, such reservations are approved by staff of the respective computing centre before being committed to the HPC system.

Advanced reservation and co-allocation is not available on other PRACE Tier-1 resources. However the larger total number of resources and technologies provides end-users with additional alternatives. For instance, PRACE Tier-1 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.

4.4.3 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.

4.4.4 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. Currently, all five HPC centers in PL-Grid support advance reservation and co-allocation of computing resources.

The main actors involved in the resource allocation process are:

- VOs and users groups use assigned resources in computational experiments.
- Resource providers 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.

4.5 Assessment of user support policies

User support policies define interfaces and type of support services end-users are provided with by an e-Infrastructure. User support covers a multitude of areas starting with definition of service access model and creation of credentials and accounts on e-Infrastructure resources through to resolution of technical issues and enabling and optimization of domain specific software.

European e-Infrastructures offer end-user support as a part of their service portfolio. The level of support services offered to end-users depends on user requirements, available expertise and allocated budget. Furthermore VOs and virtual communities may offer additional support in their areas of expertise. As such the magnitude of provided support differs between e-Infrastructures.

Further aspect to be concidered is end-user education and training. Structured and refined trainings and documentation material are essential for day-to-day support of end-users. Training sessions help users to aquire fundamental knowledge with e-Infrastructure services and share their experience with technical experts. User documentation and tutorial guide

end-users in their daily activities and answer questions that would otherwise be directed to e-Infrastructure support units.

4.5.1 PRACE Tier-0

PRACE Tier-0 operates highly complex HPC resources those operation and usage require indepth domain knowledge and expertise that is often not available in the scientific community. To compensate for the lack of end-user experience with state of the art computing systems PRACE offers a wide range of training and support services covering scientific, application and service domains.

PRACE Tier-0 offers three resource access models: early access, regular access and preparatory access. The first two models are meant to give regualr resource access and follow a review process described above. The latter, third, model offers end-users to gain limited access to Tier-0 resources for application testing, debugging and optimisation. Preparatory access grants research groups a unique opportunity to prepare scientific codes for execution on state of the art HPC resources. As a part of this access model PRACE experts supervise end-user activity and offer support in development, optimisation and deplomen of scientific applications.

PRACE organises annual training sessions to answer interests and needs of end-users. Training sessions and exersises cover a variety of topics in hardware, software and service areas helping end-users to aquire knowledge and experience in HPC domain. Furthermore these sessions provide a unique opportunity for end-users and PRACE experts to exchange oppinions and experience in working with HPC systems. Selected PRACE partners organise local training sessions that cover particular topics in greater details and help to increase community outreach.

PRACE provides users with with a reach collection of documentation and offline training material. User documentation covers all aspects of the e-Infrastructure and is meant to answer the majority of questions arising from end-users. Documentation is kept up-to-date at all times and reflects all changes in the PRACE Tier-0 infrastructure. PRACE guides and tutorials are designed to help end-users with services and tools offered in the e-Infrastructure. These documents focus on particular topics and assist end-users in mastering features they are interested in.

PRACE user help desk acts as a common access point for all infrastructure related questions. Problems and question can be submitted via a web or an email interface depending on users' preferences. The help desk is monitored at all times, so that the initial feedback can be provided within a few hours. Selected Tier-0 partners, in addition, open their local help desk systems to end-users thus increasing the overall number of available options.

Furthermore end-users are offered support in application deployment and optimization. PRACE software experts are available at every stage of a project to assist users in situations when scientific codes underperform or exhibit other unexpected behaviour.

4.5.2 PRACE Tier-1

User support services offered by PRACE Tier-1 are mostly shared with those offered by PRACE Tier-0. As mentioned in the previous section, Tier-1 offers a great range of support services, including training, documentation, help desk and application enabling. Tier-1 resources offer a richer selection of tools and services to support usage scenarious extending beyond PRACE e-Infrastructure. As such, Tier-1 pays greater attention to individual support of end-users who require assistance or experience difficulties using available system software and tools.

4.5.3 EGI

User support in the EGI e-Infrastructure is centered around application teams and support units. Application teams are responsible for development, operation and support of particular tools and services deployed in the e-Infrastructure. These can be services for access to computing and storage resources, collection of accounting data and management of virtual organisations. Each application team is tasked with providing necessary documentation, including administrator and user guides, sample use cases and tutorials, that is distributed to public. Service documentation is meant to efficiently guide end-users through access and operation of EGI services and is designed to answer questions that might arise in a conventional environment. Specific questions and incidents are handled by EGI support units that, if necessary, refer to respective application teams.

EGI support units insure sustained operation of e-Infrastructure resources and services. All major EGI services are allocated a support unit that processes enqueries received via email or through EGI GGUS help desk interface. The help desk solution employed by EGI isolates multiple support units and allows efficient management and handling of incoming requests.

Supplimentary support units are operated by EGI VOs that provide VO specific applications or services. By using EGI GGUS mechanisms VOs are able to support their users via a common help desk interface. The type of support extends to information requests, incidents, service inquiries and application support. Furthermore EGI provides project partners with an integrated XGUS help desk solution that facilitates setup and operation of an e-Infrastructure user support system.

Other EGI support activities comprise user training sessions and annual conferences organised across Europe and intended to offer user communities an opportunity to receive advance training and exchange their knowledge and experience.

4.5.4 PL-Grid

Contributes to EGI support.

5 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 below.

5.1 Deployment of 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, low level libraries and applications, without disturbing production or affecting systems performance in any undesired way. Later on MAPPER components are thoroughly tested to ensure their interoperability with each other.

To support the software validation process a pre-production infrastructure has been created. The infrastructure 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 infrastructure are 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 are passed to WP6 for deployment on the MAPPER research infrastructure. Each component is assigned a WP6 service expert responsible for co-ordinating its installation, configuration and maintenance. To facilitate the deployment process WP6 developed MAPPER Service Description Form. The form is used to describe production ready software components and provide their essential characteristics including system and security requirements. The form is to be filled out for each of the components that are deployed and provided to the responsible application expert who relies on this information during the initial installation phase.

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

MAPPER application and services are deployed on production EGI and PRACE resources after passing necessary validation and testing in the MAPPER pre-prodcution environment. To ensure successful deployment EGI and PRACE operation staff is provided with erquired supporting information including installation and administration guides, service description forms and contacts of the responsible service experts. Additionaly MAPPER partners provide necessary support for optimal service operation.

Deployment of MAPPER services on production e-Infrastructure resources is supervised by the EGI-MAPPER-PRACE task force. The task force brings together experts representing three projects and is meant as an entity for coordination and support of inter-project activities.

5.2 Monitoring and quality assurance

All deployed MAPPER infrastructure components are monitored to ensure their availability and correct functionality. Monitoring is 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 pays especial attention to monitoring solutions used by EGI and PRACE to guarantee full compatibility with the European e-Infrastructures in this area. Recommendations of the Technical Board are presented to WP6, which coordinates 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 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 is 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 is reported directly to WP6. Each incident is classified based on its impact on the MAPPER e-Infrastructure and is assigned to the service expert responsible for the failing component. The expert analyses the incident and attempts to resolve the related problem. If successful, all parties affected by the problem are informed and the solution is documented in the MAPPER knowledge base. In situations when members of WP6 cannot resolve a problem, WP4 and WP5 are asked for assistance. If necessary, the problem is escalated to the developers of the respective component.

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

5.3 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.

Detailed information on MAPPER VO and user support policies can be found in deliverable D6.3.

5.4 Integrating MAPPER services into EGI

The current status (September 2012) of integration of MAPPER operational services into EGI is summarized below:

EGI services	MAPPER/QCG
Management	Completed
Monitoring	Completed (waiting SAM release Update-19)
Accounting	In progress (waiting for a new APEL client – EMI

	3 release)
Support	In progress (requested new GGUS support unit for QCG). An independent xGUS instance will be provided for broader MAPPER support activities.
Dashboard	Not started (waiting SAM release Update-19)

5.5 Integrating MAPPER services into PRACE

6 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 is 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.

7 Acronyms and References

7.1 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
GÉANT	Gigabit European Advanced Network Technology
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

7.2 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/