SEVENTH FRAMEWORK PROGRAMME
Research Infrastructures

INFRA-2010-2.3.1 – First Implementation Phase of the European High Performance Computing (HPC) service PRACE

PRACE-1IP

PRACE First Implementation Project

Grant Agreement Number: RI-261557

D4.3.1
Cross National Programme for Access to Tier-1 Resources

Final

Version: 1.0
Author(s): Alison Kennedy, EPCC and Damien Lecarpentier, CSC
Date: 24.05.2011

Project and Deliverable Information Sheet

PRACE Project	Project Ref. №: RI-261557	
	Project Title: PRACE First Implementation Project	
	Project Web Site: http://www.prace-project.eu	
	Deliverable ID: < D4.3.14.3.1>	
	Deliverable Nature: < Report >	
	Deliverable Level: PU *	Contractual Date of Delivery: 31 / May / 2011
		Actual Date of Delivery: 31 / May / 2011
	EC Project Officer: Bernhard Fabianek	

* - The dissemination level are indicated as follows: **PU** – Public, **PP** – Restricted to other participants (including the Commission Services), **RE** – Restricted to a group specified by the consortium (including the Commission Services). **CO** – Confidential, only for members of the consortium (including the Commission Services).

Document Control Sheet

Document	Title: <Cross National Programme for Access to Tier-1 Resources>	
	ID: <D4.3.1>	
	Version: 1.0	Status: Draft
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2007	
	File(s): D4.3.1.docx	
Authorship	Written by:	Alison Kennedy, EPCC and Damien Lecarpentier, CSC
	Contributors:	Ana Bela Dias, NCF, Michael Browne, ICHEC, Giovanni Erbacci, CINECA, Georgi Prangov, NCSA, Nelly Stoyanova, NCSA, Jules Wolfrat, SARA, Stéphane Requena, GENCI, Pedro Alberto, UC-LCA
	Reviewed by:	Thomas Eickerman, JSC, David Vicente, BSC
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	DD/Month/YYYY	Draft	
0.9	11/05/011	Draft submitted for review	
1.0	24/05/2011	Final version	

Document Keywords

Keywords:	PRACE, HPC, Research Infrastructure

Table of Contents

List of Figures

List of Tables

References and Applicable Documents

[1] http://www.prace-project.eu

List of Acronyms and Abbreviations

AISBL	Association Internationale Sans But Lucratif, legal entity of the European HPC Research Infrastructure
BAdW	Bayerischen Akademie der Wissenschaften (Germany)
BoD	PRACE Board of Directors
BSC	Barcelona Supercomputing Center (Spain)
CEA	Commissariat à l'Energie Atomique (represented in PRACE by GENCI, France)
CERN	Conseil Européen pour la Recherche Nucléaire (European Organisation for Nuclear Research), situated in Switzerland
CINECA	Consorzio Interuniversitario, the largest Italian computing centre (Italy)
CINES	Centre Informatique National de l'Enseignement Supérieur (represented in PRACE by GENCI, France)
CoI	Co-Investigator (major collaborator on a joint proposal)
CPU	Central Processing Unit
CSC	Finnish IT Centre for Science (Finland)
CSCS	The Swiss National Supercomputing Centre (represented in PRACE by ETHZ, Switzerland)
DECI	Formerly DEISA Extreme Computing Initiative, now Distributed European Computing Initiative
DEISA	Distributed European Infrastructure for Supercomputing Applications. EU project by leading national HPC centres.
EC	European Community
EPCC	Edinburgh Parallel Computing Centre (represented in PRACE by EPSRC, United Kingdom)
EPSRC	The Engineering and Physical Sciences Research Council (United Kingdom)
ETHZ	Eidgenössische Technische Hochschule Zuerich, ETH Zurich (Switzerland)
ESFRI	European Strategy Forum on Research Infrastructures; created roadmap for pan-European Research Infrastructure.

FZJ	Forschungszentrum Jülich (Germany)
GCS	Gauss Centre for Supercomputing (Germany)
GÉANT	Collaboration between National Research and Education Networks to build a multi-gigabit pan-European network, managed by DANTE. GÉANT2 is the follow-up as of 2004.
GENCI	Grand Equipement National de Calcul Intensif (France)
HET	High Performance Computing in Europe Taskforce. Taskforce by representatives from European HPC community to shape the European HPC Research Infrastructure. Produced the scientific case and valuable groundwork for the PRACE project.
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing
IDRIS	Institut du Développement et des Ressources en Informatique Scientifique (represented in PRACE by GENCI, France)
JSC	Jülich Supercomputing Centre (FZJ, Germany)
KTH	Kungliga Tekniska Högskolan (represented in PRACE by SNIC, Sweden)
LINPACK	Software library for Linear Algebra
LHC	Large Hadron Collider, the world's largest high-energy particle accelerator
LRZ	Leibniz Supercomputing Centre (Garching, Germany)
MoU	Memorandum of Understanding.
MPI	Message Passing Interface
NCF	Netherlands Computing Facilities (Netherlands)
NDGF	Nordic Data Grid Facility
Open MP	Open Multi-Processing
PFlop/s	Peta (= 10^{15}) Floating-point operations (usually in 64-bit, i.e. DP) per second, also TF/s
PI	Principal Investigator (lead scientist on a proposal)
PRACE	Partnership for Advanced Computing in Europe; Project Acronym
PRACE-PP	PRACE Preparatory Project
PRACE-1IP	PRACE First Implementational Phase Project
PRACE-2IP	PRACE Second Implementational Phase Project
PSNC	Poznan Supercomputing and Networking Centre (Poland)
RI	Research Infrastructure
RP	Resource Provider (TeraGrid initiative)
RT	The Nordic Resource Trading Project
SARA	Stichting Academisch Rekencentrum Amsterdam (Netherlands)
SNIC	Swedish National Infrastructure for Computing (Sweden)
STFC	Science and Technology Facilities Council (represented in PRACE by EPSRC, United Kingdom)
TFlop/s	Tera (= 10^{12}) Floating-point operations (usually in 64-bit, i.e. DP) per second, also TF/s
Tier-0	Denotes the apex of a conceptual pyramid of HPC systems. In this context the Supercomputing Research Infrastructure would host the Tier-0 systems; national or topical HPC centres would constitute Tier-1, regional of local centres will form the Tier-2 level
TRAC	TeraGrid Resource Allocations Committee
UC-LCA	University of Coimbra, Laboratory for Advanced Computing (Portugal)
VO	Virtual Organisation, an alliance of entities pooling and sharing (computational) resources electronically for the purposes of furthering their research

Executive Summary

The purpose of this document is to provide input and guidance to the PRACE AISBL (the legal entity of the European HPC Research Infrastructure) on the selection of a model of resource exchange to be used by PRACE in cross-national programmes for access to nationally-operated HPC facilities. The model chosen will initially be used to provide a framework for exchange of computing resources between those PRACE Tier-1 sites who are contributing nationally-funded CPU cycles to the PRACE-2IP project for allocation via a peer review process such as e.g. the single project access scheme DECI (formerly operated by DEISA and soon to be operated by PRACE).

However, PRACE's longer term aim is to develop a resource exchange model which is capable of extension to provide a framework for resource exchange between different Tier levels in the HPC ecosystem (e.g. between Tier-0 and Tier-1 or between Tier-1 and Tier-2) or between European resources and those of other continents or nations (e.g. between European and US projects) or between CPU and other types of resources which are essential to computational scientists such as data storage, visualisation resources, etc. This ensures the potential for adaptability to meet the evolving requirements of computational scientists in the European Research Area.

To date, the needs of the Tier-1 HPC user community in Europe have been largely supported by national activities and latterly by the EC Framework Programme-funded DEISA projects. The PRACE-2IP project will integrate some of the DEISA project activities into PRACE to provide an integrated European HPC ecosystem with seamless access to capability computing resources covering various Tier levels. In practical terms, therefore, the initial objective of this task is to support the integration process by analyzing the existing DECI resource exchange model against the PRACE Tier-1 resource exchange requirements. It can subsequently recommend any changes or enhancements necessary to facilitate the migration of DECI from DEISA to PRACE and to support the allocation decision process of the first DECI call in PRACE. The PRACE resource exchange model, however, will not necessarily be very closely based on the DEISA DECI model, but will analyze medium-term European requirements and study international best practice, before recommending a solution to the PRACE AISBL.

This document specifically focuses on models for CPU resource exchange between PRACE partners providing Tier-1 resources. Subsequent deliverables in PRACE-2IP will look at extending the model beyond the original concept.

1. Introduction

The objective of Task 4.3 of the PRACE-1IP project is to develop and pilot practices to stimulate Pan-European HPC access. This includes the development and evaluation of a model for access to Tier-1 HPC resources based on an interchange of cycles among Tier-1 computing centres of different countries. In undertaking this task, PRACE will take account of the experiences and good practices developed by the DEISA and DEISA2 projects through the DEISA extreme computing initiative (DECI, which is continuing in PRACE) and by, amongst others, HPC-Europa and HPC-EUROPA2, the European Commission's Transnational Access and Mobility programme in HPC.

When this task was originally proposed, the deliverable was intended primarily to provide information on which to make informed decisions about a future, but as yet undefined, Tier-1 resource exchange programme. However, the recently funded PRACE-2IP project has the stated objective of integration of Tier-1 resources, services and user communities within the PRACE RI with a specific focus on re-using the successful components of the DEISA projects. DECI calls (DEISA Extreme Computing Initiative), the current single-project cross-national access scheme for Tier-1, will be continued with the same acronym but now meaning Distributed European Computing Initiative.

A key feature of PRACE-2IP will be its focus on users and on ensuring that the existing community of PRACE Tier-0 users is expanded and enhanced by the addition of new users both from the former DEISA (and then PRACE) Tier-1 community and entirely new users who are out-growing regional or national resources. By supporting an ever-growing community of users, the future growth of HPC use in Europe by science and industry will be assured. Some of them will continue to have their needs met by Tier-1 access, but others will progress from Tier-1 to Tier-0, fully supported by PRACE's application enabling and optimisation activities. A PRACE resource exchange model is therefore of importance to promote cross-national usage of Tier-1 systems but also to facilitate and stimulate migration to Tier-0 and to support the development of cross-national and interdisciplinary research collaborations with multiple computing needs.

PRACE-2IP, which will start in September 2011, has a workpackage which focuses on the integration of the resources and services formerly provided by DEISA into the European HPC ecosystem coordinated by PRACE AISBL. It will build on the findings of this deliverable to develop a fully functioning cross-national Tier-1 resource interchange system. Users will ultimately see a consistent European HPC ecosystem that addresses both the access to Tier-0 and to Tier-1 resources and the migration of applications from one performance level to the next.

This document therefore presents resource interchange concepts to PRACE members in general and the PRACE AISBL and Council, in particular.

Concepts will subsequently be tested by means of pilot cases with a view to defining and finalising a model of resource exchange to be developed by PRACE-2IP and finally be used by PRACE AISBL. This model will primarily be used as a test case to exchange HPC resources (CPU) between Tier-1 centres participating in the PRACE-2IP project, but the intention is to define a model that can ultimately be broadened to include:

- Exchange of resources between PRACE AISBL and other infrastructures e.g. TeraGrid and its successor projects
- Exchange of different types of resources e.g. data storage

- Exchange of resources between different levels of the HPC ecosystem (where the key differentiator for the purposes of this task is the infrastructure funding model and governance (PRACE versus national or topical), rather than the absolute or relative performance of the HPC system)

The current resource exchange model used by DECI in DEISA will form the basis of the initial pilot to facilitate the earliest possible implementation of DECI during PRACE-2IP.

Following the evaluation of the DECI pilot and of a second pilot (looking at extending the scope of resource exchange to include additional factors such as data storage), a follow-up deliverable will be produced in month 24, which will document all aspects of the PRACE resource exchange model to be approved by the Council of PRACE AISBL.

2. Methodology

The first question which the workpackage addressed was related to the purpose of the intra-national access mechanisms. It was agreed that the main purpose was to determine a fair exchange rate mechanism by devising a model which allowed the value of individual contributions of CPU resources to be measured and valued in an objective and transparent way. We assumed that this would be done by means of an agreed exchange rate based on a common unit of currency so that resources could easily be exchanged between any two partners in the project (rather than being exchanged, by barter, through bi-lateral agreements between partners).

For the purposes of the initial exchange rate model, there is an initial assumption that PRACE partners' contributions and usage are expected to balance over time i.e. that no partners are expected to contribute more than they use on an on-going basis, the concept of juste retour. If that were not the case, then we would need to extend the model to assess the monetary value of contributions instead of just the relative value. The model will also assume that the resources being exchanged are free at the point of use for researchers (i.e. it excludes the sale of CPU resources to third parties).

It was decided that there should be an agreed threshold for contributions (relative or absolute) in terms of machine size, queue size and configuration, and amount of core-hours (i.e. a minimum specification for a PRACE Tier-1 system which is included in the resource exchange model). This would be determined initially by the Management Board, and at a later stage approved by PRACE AISBL on a call-by-call basis.

It was then agreed that, in order to develop a suitable resource exchange model, the following actions would be undertaken, to provide a framework for evaluation of possible solutions and to further define requirements.

- Produce a working definition of Tier-1 for the purposes of this task (and agree the definition with the PRACE AISBL Council) (See Section 3)
- Undertake a review of existing models for resource exchange, collecting information from other projects and initiatives (See Section 4)
- Analysis of supplementary information, collected by project-based user questionnaires, to provide additional input on user requirements for cross-national access programmes (See Section 5)

This deliverable documents the outcomes of each of the above actions and describes their contribution to the development of a proposed model for resource exchange.

A further four potential actions were identified to be undertaken at a future date, in order to provide supplementary information for the recommendation of a model to be approved and adopted by PRACE AISBL and ultimately for the subsequent deliverable in M24 which will document the model of resource exchange to be adopted for PRACE AISBL:

1. Agree a framework for evaluating existing resource exchange models
2. Agree the requirements for the PRACE AISBL resource exchange model (in collaboration with the Board of Directors and the Council of PRACE AISBL) – combine input from the review of existing models, the questionnaires and the evaluation model used during the pilot calls for Tier-1.
3. Analysis of supply and demand for HPC resources for Tier-1 and large Tier-2 resources within each country (to help to forecast future demand and to supplement information collected via the pilot call and historical data from DEISA)
4. Examine the potential of a resource bank to manage supply and demand and facilitate resource exchange

Parameters for the PRACE Resource Exchange Model

A second important consideration was to discuss the possible mechanisms for determining the actual exchange rate to be used in the PRACE resource exchange model for Tier-1 and to weigh up the relative costs and benefits of each of these approaches.

There were three basic options (although combinations of one or more methods are also likely).

- Subjective e.g. by expert agreement of the service providers via a specially convened committee
- Objective e.g. by measured performance (e.g using Linpack or PRACE applications benchmarks)
- Monetary e.g. by monetary value of partner contributions (additional cost, market cost etc.)

Whichever method or combination of methods is chosen, PRACE AISBL needs also to decide whether all contributions are of equal value to users or whether we should potentially offer incentives to modify, encourage or reward specific user behaviour or to reward partners who provide a feature-rich HPC environment.

The task also needs to determine the threshold for considering a revision of the exchange rate (i.e. should there be a tolerance or should every change/upgrade to a system trigger a review of the exchange rate or should exchange rates be fixed for a finite period of time and if so for how long?)

Rather than starting from scratch, it was decided to look at what other projects have done previously in similar circumstances and to evaluate the suitability and applicability of their approach to PRACE. In particular, our approach involves:

- Looking at exchange rate models and mechanisms in use elsewhere
  - In HPC and Grid (Europe/international)
  - In other infrastructures (just Europe)
- Assessing the strengths and weaknesses of these models and appropriateness for PRACE AISBL

The final stage involves identifying some pilot cases, running pilots to test the concepts and evaluating the results. The first pilot will involve assessing the extent to which the DEISA resource exchange model meets the needs of PRACE AISBL, with particular reference to the running of DECI calls in the framework of PRACE AISBL.

3. Tier-1 Definition

To quote from the Description of Work, "The work done at the Tier-1 level has a major impact on Tier-0 and high quality support from Tier-1 level services is required to provide a set of scalable applications and optimal load for Tier-0 utilisation. For that reason it is of utmost importance for the PRACE RI to have the support of a good-working Tier-1 infrastructure, with sufficient interoperability with Tier-0."

The first step in the process of developing a model for interchange of cycles among Tier-1 centres of different countries is to define (for the purposes of this activity) what we mean by a Tier-1 system and Tier-1 centre. These definitions will initially be used in WP4, but may be expanded or refined for use in PRACE-2IP and finally in PRACE AISBL.

PRACE already has a working definition of a Tier-0 system, drawn up during the PRACE-PP project. This definition is quite short and is as follows:

- the cycles of the system are allocated through the PRACE peer review process (in other words: the governance over the system is in the hands of PRACE)
- the system is part of a 100 Mio Euro installation, and is of capability-class (meaning that jobs can use the full machine)

WP4 therefore set out to provide a suitable working definition of Tier-1, later approved by the project Management Board as such.

The definition of Tier-1 adopted by WP4 should, as far as possible, be analogous to this definition. However, following discussions between the partners in WP4, it became obvious that, for the purposes of our task, we needed a definition of what services a Tier-1 centre needed to provide, as well as a definition of what constituted a Tier-1 system. These definitions will initially be used in WP4, but may be expanded or refined for more general use.

In order to participate in the PRACE Tier-1 resource exchange model, partners must be able to offer resources which meet the Tier-1 system criteria and an operational model which meets the Tier-1 centre criteria. (Note that we expect a Tier-1 system to be a single resource i.e. not a grid, but that the services provided by the Tier-1 centre may be distributed, if this is the usual way that HPC services are provided in that country). Partners must also provide a **Tier-1 contribution** of sufficient size to make their participation in a resource exchange scheme viable.

Tier-1 centre:

- offers user training in HPC programming;
- offers technical and user support via a help desk;
- offers applications support (code development and optimization);
- offers advanced development platforms for HPC (i.e. pre-cursors of future Tier-0 architectures);
- offers services to support Grand Challenges and provide a ramp to allow scientists to move to Tier-0 systems;

- is willing to implement necessary middleware for integrating the site into the PRACE Tier-1 infrastructure;
- is willing to sign a SLA with PRACE (i.e. to address a list of pre-requisite requirements).

Tier-1 system:

- cycles or science projects allocated via external (national) peer review;
- supports the work of researchers nationally and/or within a region of Europe (if such arrangements are in place from the Tier-1 centre);
- has a hardware and software configuration able to meet the computational needs of a range of science areas;
- has a high enough performance to meet most of the national requirements below Tier-0 (will typically be no more than one order of magnitude lower than the Tier-0 systems);
- serves as a development platform for the most scalable codes (which will be directed towards the Tier-0 systems);
- has batch queues configured to support capability computing (use of maximum available computing power to solve a problem in the shortest amount of time) although the system may support both capability and capacity computing;
- is connected to the PRACE infrastructure with sufficient network/bandwidth to facilitate the exchange of large datasets (with other Tier-1 and Tier-0 systems) and participates in PRACE Tier-1 resource exchange.

Tier-1 Contribution:

- usually 5% (or more) of the resources available on a Tier-1 system annually (the typical DEISA contribution);
- resources committed should be sufficient for at least two DECI projects per year (typical DEISA contribution was for ~5 DECI projects per year);
- for smaller centres where 5% of the total resources is too small to support two DECI projects per call, there is the possibility of banking resources and contributing a larger amount of resources to alternate calls e.g. 10% to every second call (the feasibility of this approach will be investigated during the resource exchange pilots).

4. Review of Existing Models

In this section, several existing resource exchange models are briefly reviewed and differences in user resource provisions are identified in order to define the way towards the development of a PRACE resource exchange model. The model adopted must be capable of evolution, both as Tier-1 needs and expectations change but also to accommodate additional PRACE partners and a wider range of resources to be exchanged. In other words, it must be able to scale to a large number of resource providers and to a range of different resource types.

The main goal is to develop and evaluate a PRACE model for access to Tier-1 HPC resources based on an exchange of cycles among Tier-1 computing centres of different countries.

This section describes how input to the decision process was collected, based on the experiences of different projects and initiatives by:

- Contacting a number of e-Infrastructures initiatives and projects and asking for information on resource exchange models using some pre-defined procedures and a template (to ensure that comparable information was collected from each project);
- Analysing the strengths and weaknesses of identified models (as potential models for PRACE);
- Providing input both strategical and operational to the process of defining the PRACE resource exchange model based on best practice from other projects.

The overall objective is to support the evaluation of the existing resource exchange models, identifying suitable organisational and technical guidelines, building blocks and practices that help in the overall formulation of the PRACE resource exchange model. The framework includes an analysis of a number of e-Infrastructure projects, presenting clearly which existing resource exchange models are most appropriate for our requirements and might be taken into consideration in the process of implementing the PRACE resource exchange model. As a starting point, the DEISA model was examined, but it was felt to be extremely important that the analysis was spread as broadly as possible to investigate various other models and specifications which directly or / and indirectly could also act as reference models for PRACE.

In undertaking the analysis, an important precondition was set – that the main criteria on which basis the identified e-Infrastructures initiatives and projects will be compared should be determined in advance. In addition, this section's approach excludes the assumption that PRACE resource exchange model should be built from scratch, dedicating extra effort and time.

The aim is to provide some recommendations about the partial or full integration of aspects of the reviewed models into the PRACE Tier-1 model, bearing in mind the specific PRACE environment.

4.1 Data collection process and criteria comparison of existing resource exchange models

This subsection elaborates and describes the mechanism for data collection as well as the main criteria for comparison between existing resource exchange models in order to define the most appropriate characteristics and general practices that can be applied to the PRACE Tier-1 resource exchange model. It can be viewed as a step in the process of shaping a framework that allows us, in an objective and transparent way, to present the overall analytical process from demonstrating how we are gathering information through to the way we are measuring the collected information up to the following implementation.

Data collection process

Initially, a data collection framework that standardised and facilitated information collection process was set up. It was based on clearly defined criteria balancing between HPC and Grid resources exchange models on one side and national, regional, European and global dimensions of the reviewed models on the other side. The main criteria included: collecting brief but relevant information; documenting who is involved; the size and range of HPC facilities involved; how long it has been operating for; how widely used it is; and perceived strengths and weaknesses.

As a result of the defined criteria and follow-up discussions, a tight template form was prepared for collecting information about existing resource exchange models. Additionally, the geographical allocation of the reviewed existing resource exchange initiatives was noted and projects grouped as following:

- National initiatives and project: an example is US Teragrid;
- Regional[1] European initiatives and projects;
- European initiatives and projects;
- Global initiatives and projects;
- Other initiatives and projects.

At the next stage, the template form was distributed for information collection asking interviewers to concentrate on HPC and Grid resources, but if other resources are exchanged, they can collect information (which will be of use also for the second deliverable).

Information about the PRACE Tier-0 peer-review process, although already familiar to the PRACE Council, is included for comparison purposes.

Criteria for benchmarking of existing resource exchange models

As a next step, after defining the data collection process, we needed to elaborate a set of broad criteria for benchmarking of existing resource exchange models for which information had been obtained. For the purposes of this report, the following criteria are defined as benchmarks:

- Level of maturity – implementation stage and track record of the project at any point of time;
- European dimension – national, regional, European and global;
- Relevance to PRACE resource exchange needs and requirements – exchange of HPC resources – CPU; data storage etc.;
- Recognition by European HPC community – visibility from science and engineering communities for supporting and solving certain scientific challenges;
- Governance migration to PRACE environment – ease of process and procedural integration in PRACE environment mitigating operational and administrative disruption.

The identified criteria are not intended to be comprehensive or used for in depth analysis of the collected resource exchange models, but they need to provide overall insights and impressions about the potential replication or partial integration of other existing models in the forthcoming PRACE Tier-1 resource exchange model.

4.2 Analysis of e-Infrastructure resource exchange models - current practices and future development

For this task, we compiled a short list of 7 projects to contact to discuss resource exchange models, as the list was restricted to HPC or Grid infrastructure projects. Looking by geographical distribution point of view, information was gathered for 4 projects with pure European dimension, 2 projects combined the European efforts with countries outside Europe and 1 project was nationally defined coming by a country outside Europe – USA. The following projects were identified and contacted:

[1] Here, and throughout the document, region (or regional) is used in the sense of a region in Europe, and not a region within a country.

Name	Type	Geographical Region
DEISA2	HPC Infrastructure	Europe
TeraGrid	HPC Infrastructure	USA
EGI	Grid Infrastructure	Global
HPC-Europa2	Transnational Access to HPC (includes researcher mobility)	Europe
Nordic Data Grid Facility	Regional Grid	Nordic countries
HP-SEE	Regional HPC Initiative	Southeast Europe
LinkSCEEM	Regional HPC Initiative	Cyprus and Mediterranean rim countries

Table 1 List of projects contacted to ask about resource exchange

The main efforts were concentrated on the collection of as many as possible existing resource exchange practices going beyond the DEISA model. It should allow us to verify or reject the initial presumption that PRACE Tier-1 resource exchange framework should replicate the DEISA DECI initiative on a broader base. In this context, the follow up analysis will review, first and foremost, the DEISA resource exchange model comparing it to other identified e-Infrastructure initiatives and projects using the defined criteria in Section 4.1 Data Collection process and criteria comparison of existing resource exchange models To do this, every resource exchange model is briefly described providing answers to questions such as: what kinds of resources are being exchanged? What are the mechanisms for exchanging these resources? etc.

Once again, it is extremely important to underline the fact that the presented projects comparison is only for the purposes of defining PRACE resource exchange model elements and conditions under which PRACE AISBL can offer its users effective, efficient and quality Tier-1 services. The benchmarking criteria cannot be used for individual project evaluation as well, as this is not the intention of this document.

PRACE Tier-0

The Partnership for Advanced Computing in Europe (PRACE) is a persistent pan-European Research Infrastructure for High Performance Computing (HPC). PRACE forms the top level (Tier-0) of the European HPC ecosystem and provides access to computing resources at the highest performance level as well as sophisticated services for scientific simulations in all fields of research and engineering. PRACE Tier-0 is a distributed Research Infrastructure with supercomputer systems of Petaflop level installed in various PRACE member countries. PRACE has at present 2 types of access: Project (allocation of resources for one year) and Preparatory (allocation of resources for code testing and scalability development on a rolling basis). Programme access (allocation of resources for multiple years) is also foreseen for the near future. Researchers can centrally apply to the Tier-0 resources available through calls. For Project Access there are two annual calls and the call for Preparatory Access is open all year round. Researchers can access the PRACE Tier-0 resources through a single European peer review process based on scientific excellence and the proven need of Tier-0 resources for execution of the project. All proposals undergo PRACE peer reviewing. Project proposals are technical and scientifically assessed by independent reviewers. Technical assessment regards

mainly evaluation of feasibility of the project and optimal usage of the computing resources requested. Scientific assessment concerns scientific quality, innovation and impact of the project and is performed by at least 3 independent reviewers with expertise in the scientific field of the proposal. Reviewing is anonymous, i.e. the names of the reviewers are kept anonymous and conflicts of interests are identified. Applicants have the right to reply to the comments of the technical and scientific reviewers. Proposals are ranked by the Prioritisation Panel on the basis of the reviewers' reports and the applicants' replies. Resources are allocated on the basis of a single ranking list and resources are allocated until being exhausted or the remaining proposals are below the quality threshold defined by the Prioritisation Panel. Proposals for Preparatory Access undergo technical assessment. PRACE AISBL has a Scientific Steering Committee consisting of internationally recognized scientists from various scientific fields who advise PRACE AISBL on all matters regarding scientific usage of the PRACE AISBL resources, including the methodology used in the peer review process.

Advantages

- Good level of maturity, designed to to take account of international best practice in peer-review, based on input from many national peer-review experts
- Transparent procedures, understood and accepted by PRACE partners and Tier-0 user community

Challenges

- A single peer review panel is almost certainly not feasible for Tier-1 calls, given the large number of proposals expected. The PRACE Tier-0 model may struggle with volume as the number of Tier-0 systems grows, so it does not seem likely that Tier-1 and Tier-0 could simultaneously be reviewed by the same panel and mechanism.
- Tier-0 applicants apply to one particular system, therefore resource exchange between Tier-0 systems has not been applicable to date and resources are not exchanged per se
- In future, applicants may wish to apply for both Tier-1 and Tier-0 access e.g. for multi-scale applications or projects running more than one code with different performance characteristics. A close degree of co-operation in resource allocation and exchange between the two tiers would facilitate this.

Distributed European Infrastructure for Supercomputing Applications (DEISA)

DEISA is a consortium of leading national Supercomputing centres that aims at fostering pan-European world-leading computational science research. The project is exchanging CPU cycles between thirteen national supercomputer centres so that researchers from any European country can obtain peer-reviewed access to the most suitable architectures and machine configurations for their science. The local CPU hours are converted to DEISA processor hours and all allocations are made in DEISA units. The conversion rates are set in advance of allocations being made. Applicants describe their computational requirements and make an estimate of the resources required in whatever units of CPU they are most familiar. The proposals undergo both a technical review (from experts within the centres who are partners in the project) and a scientific peer-review from external experts. The technical review assesses their suitability for the DEISA infrastructure in general and in particular, for the particular architecture and resources requested. Scientific peer-review is undertaken by a national peer-review committee, based on the criteria set up by DEISA (and publicised to applicants). Reviewers are anonymous. Each peer-review committee manages conflict of interest according to their national procedures. Projects are then ranked. Based on these reviews and their rankings, taking into account user expressed preferences for particular

architectures or machines wherever possible, projects are allocated to machines by an allocation committee. Awards can be scaled back from the amount requested in the proposal by the allocation committee. Awards are made in DEISA core-hours, but users are told how this relates to their original request.

What are the strengths and weaknesses of DEISA as a case example to future PRACE resource exchange model? Here some of the main advantages and potential challenges will be pointed out:

Advantages:

- High level of maturity - continues to develop and support the pan-European distributed high performance computing infrastructure established since 2002 within the predecessor project DEISA that was funded under FP6. During the two DEISA projects there were organised 6 DECI Calls duly providing information about awarded projects;
- A clear pan-European approach – since its establishment, DEISA has deployed and operated a persistent, production quality, distributed supercomputing environment with continental scope. By extending the European collaborative environment in the area of supercomputing, DEISA has paved the way towards the deployment and operation of a persistent cooperative European HPC ecosystem, as suggested by ESFRI, and provided a turnkey operational solution for it.
- Strong relevance to PRACE resource exchange needs and requirements - the project is exchanging CPU cycles (the unit of exchange is the DEISA processor core hour) among national supercomputer centres, so that researchers from any European country can obtain peer-reviewed access to the most suitable architectures and machine configurations for their science. Local CPU hours are converted to DEISA processor hours and all allocations are made in DEISA units. The DEISA benchmark suite is run on each new system or upgraded system to provide a measure of achievable applications performance with a range of user codes. The results of the benchmark runs provide input to the setting of the conversion rates. The conversion rates are set in advance of allocations being made.[2] Also support for application enabling can be requested.
- Obvious recognition by European HPC community – since 2005, when the first DECI call was announced, more than 230 proposals were received by 364 different researchers from 25 European countries have participated in the scheme to date as Principal Investigators or as Co-Investigators, along with collaborators from four other continents.
- Easy governance migration to PRACE environment – from the 14 principle and associated partners, only 2 organisations are not represented and have not participated in PRACE. Furthermore, the majority of sites that form the DEISA research infrastructure are operated by PRACE members and have a high level of coherence between their representatives participating in both initiatives.

[2] The setting of CPU conversion rates for DEISA is described further in the public DEISA2 deliverable *Maintenance of the DEISA Benchmark Suite in the First Year*

Challenges:

- PRACE represents 20 countries with reasonable expectation the number of members to grow and therefore the PRACE resource exchange model will need to integrate more and more new infrastructures. This means that PRACE Tier-1 model needs to take into consideration the forthcoming enlargement trends.
- Integration of PRACE Tier-1 resource exchange model with the overall PRACE framework. Further investigation can be conducted in the several directions including exchange of resources between both Tiers and strengthening the role of Tier-1 as an enabler of future Tier-0 projects.
- Expansion in the portfolio of PRACE Tier-1 offered services including gradual development starting by the provision of a resource bank to assessing the medium term perspective to foster research mobility such as under HPC-Europa2 project.

In order to verify the DEISA DECI model as a PRACE appropriate model additional comparison with all other existing resource exchange models need to be done.

Other existing resource exchange models – description, kinds of resources being exchanged and mechanisms for exchanging these resources

Additionally to DEISA, we contacted and gathered information from the following projects and initiatives:

1. TeraGrid – an open scientific discovery infrastructure combining leadership class resources at eleven partner sites to create an integrated, persistent computational resource. It is therefore similar in concept and objectives to DEISA. Using high-performance network connections, TeraGrid integrates high-performance computers, data resources and tools, and high-end experimental facilities around the United States of America. Currently, TeraGrid resources include more than 2 petaflops of computing capability and more than 50 petabytes of online and archival data storage, with rapid access and retrieval over high-performance networks. Researchers can also access more than 100 discipline-specific databases. TeraGrid uses a normalised unit (Service Unit). This is based on a Cray XMP core-hour. Resources used are reported in local units which are then converted to normalised units[3]. Billing rates can potentially be different. TeraGrid has one allocation process for startup and education requests and research requests. Startup and Education requests are accepted at anytime and reviewed by TeraGrid staff. Research requests are reviewed quarterly by the TeraGrid Resource Allocations Committee (TRAC). The TRAC consist of volunteer experts from the faculty and staff of U.S. universities, laboratories, and other research institutions. Each TRAC Research Request (New / Renewal / Supplement / Appeal) is reviewed by members of the TRAC. The reviewers evaluate the merits of the proposal, based on the predefined criteria. Each quarter requests are assigned to reviewers. After a period of one month to complete the reviews, the reviewers convene at a TRAC meeting to determine a Recommended Allocation for each TRAC Research Request (and any large Supplement or Justification). Allocations can span multiple years. The review includes a technical review for feasibility and a scientific assessment. For

[3] SU Conversion Calculator - Use this form to convert service units (SUs) on one platform to the equivalent amount on another platform using the accepted conversion factors. This calculator uses a weighting factor based on the performance of resources against the HPL benchmark, which is used to determine the Top 500 list of HPC systems.

applications that have received an NSF grant or equivalent financial support to fund the work-force from a national agency or foundation, the science case is not re-assessed but only the appropriateness of the computational approach. NSF supported Computational Research Plans have priority over all non-NSF supported components.

2. European Grid Initiative (EGI) – provides access to resources, i.e. logical and/or distributed entities to be shared by end-users – for example CPUs, data storage, instruments and digital libraries. Resources are contributed by Resource Centres, also known colloquially as sites. The Resource Centre is the smallest localised or geographically distributed administration domain, where EGI resources are managed and operated. Access is granted to Virtual Research Communities, by way of Virtual Organisations (VOs), and these communities have to negotiate with resource providers, in general the National Grid Infrastructures (NGIs), for the allocation of resources. Within EGI there is no model for the exchange of resources among resource providers. Preliminary or preparatory access also depends on the facilities offered by individual resource providers.
3. HPC-Europa2 – the core activity of the project involves the provision of HPC Transnational Access (TA) by seven national HPC Centres. During the four years of activity of the project (January 2009 - December 2012) the transnational activity will provide HPC resources and services, specialist support, computational tutoring and scientific collaboration to more than 1,000 European researchers, greatly contributing to create an European HPC eco-system at the service of computational sciences. The project grants CPU cycles available at the seven national supercomputer centres involved in the TA activity. The TA is granted to researchers from any European country on the basis of a peer-review selection of the applications. Each TA centre is responsible for providing users with access to HPC facilities at their centre together with local scientific and HPC support. This service is enhanced and supplemented with a co-ordinated suite of training and consultation activities. Furthermore the sites provide in kind a further amount of cycles from their national HPC systems. Around 20% of the project financed cycles are in kind contribution. The unit of exchange is the Allocation Unit (AU). The AU is defined as the "computational power delivered by a computer executing for one hour at the sustained rate of one GFlop/s", as derived by the top500 list (www.top500.org). Each researcher applies for access to a specific HPC system of one of the seven TA sites, asking for a given amount of CPU core hours on that system, sufficient to realise his scientific project. The proposals undergo both a technical evaluation and a scientific peer review. The technical evaluation is done by a member of staff from the specified TA centre, who will comment on aspects such as the suitability of the facilities requested, the availability of resources requested (such as packages, disk space, compilers and other tools etc), the amount of training likely to be required to enable the applicant to make effective use of the facilities and the feasibility of the draft work plan.
4. Nordic Data Grid Facility (NDGF) – a collaboration between the Nordic countries (Denmark, Finland, Norway, Sweden) established in 2002 to allow researchers in these countries to create and participate to large computational challenges, in particular those resulting from CERN LHC experiments. NDGF is a production grid facility that leverages existing, national computational resources and grid infrastructures. It provides and operates the Nordic Tier-1 for WLCG – it is the only distributed Tier-1 of the LHC computing grid. NDGF does not have a resource exchange system *per se* but does, however, allow access for Virtual Organizations (VO) to a set of resources (mostly computing and storage resources) distributed among its centres. The Resource Trading (RT) project (2010) examined the possibility of building a Nordic market place where the resources could come from the

participating computing centres or the individual users. The most feasible scenario was named "pooled resources". In this model, each participating country allocates a set of resources to a common pool from where the users (research groups) can select the services they wish to use. Resource exchange does not happen by direct trading between users, centres, or countries. In the proposed "pool model", resources are allocated through a continuous process done at national level. Each country can grant as much resources as they want. The national users receive then a corresponding amount of credits for accessing services from other countries.

5. HP-SEE – brings together the National HPC infrastructures in the region of South Eastern Europe and the regional Virtual Research Communities of Computational Physics, Computational Chemistry and Life Sciences, aiming at enabling those user communities to get access to HPC resources for their scientific work. The countries that participate in the project are: Greece, Bulgaria, Romania, Turkey, Hungary, Serbia, Albania, Bosnia-Herzegovina, FYR of Macedonia, Montenegro, Moldova, Armenia, Georgia, and Azerbaijan. The project aims at defining a clear exchange model between the partner countries that will be valid even after the duration of it. Input from similar initiatives is welcome. At the moment there is not a specialised unit of exchange in the project, and the CPU core hours provided by the accounting systems are used as a rough estimate of resource provision. The project has set up an Applications Review Committee (ARC) that consists of experts from the three scientific fields of strategic importance for the region (Computational Chemistry, Computational Physics, Life Sciences), infrastructure experts and the technical manager of the project. It is the responsibility of the ARC to evaluate the amount of resources that applications require, the prioritization of applications, the allocation of resources to the available systems and the approval of new applications that are not specified in the DoW of the HP-SEE project.
6. LinkSCEEM – the provision of Tier-1 supercomputing resources to regional user groups in the Eastern Mediterranean (www.linksceem.eu). The objective is to facilitate the creation of an HPC eco-system for the scientists in the region and enable them to integrate their small cluster resources with Teraflops computers coupled with the appropriate training and technical support. A resource allocation mechanism has been developed and will be applied to upcoming users; a resource exchange procedure has not been developed yet. Maybe that at a later point in time a resource exchange may be an opportunity for LinkSCEEM users. Of prime importance is the set up of a project Resource Allocation Committee administering user requests for access to partner resources. LinkSCEEM is currently implementing the preparatory access procedure. It is written around the concept of core-hours without any discrimination between architectures.

Comparison between DEISA and other existing resource exchange models

Generally, the comparison between DEISA and other existing resource exchange models, based on the criteria for benchmarking, shows that:

- Ignoring the regional dimensions of the Nordic Grid, HP-SEE and LinkSCEEM projects, it can be easily differentiated that these projects need some extra time to reach DEISA's resource exchange model level of maturity. Further investigation of these projects can be taken into consideration at later stage of PRACE resource exchange model development and deployment;
- EGI approaches the DEISA model when we measure the level of maturity, European dimension and recognition by European HPC community, but it can be predicted

greater difficulties in the governance migration to PRACE environment compared to the DEISA model and even to the HPC-Europe2 project. In addition, the applied EGI resource exchange model is restricted to VOs, which imposes some doubts about how this model would be practically reacting in one extended environment including much more participants;

– HPC-Europa2 project meets at the highest level the criteria - level of maturity, European dimension, recognition by European HPC community and governance migration to PRACE environment, but significantly deviates from the PRACE primarily goals (HPC-Europa2 set up Transnational Access model) for setting up a resource exchange model based on CPU cycles and data storage exchange between different European countries. HPC-Europa2 differs from some of the other initiatives in several respects – the amounts of EC-funded CPU on offer are comparatively small, the access period is short and is tied to a visit or at least a "virtual visit" and the main purpose of the activity is not to undertake production runs per se, but to enable researchers to establish and build collaborative links with researchers at host institutions. However, as a model with different characteristics from PRACE, designed to meet the needs of a sub-section of the HPC community, mainly young researchers who would not otherwise have access to major HPC facilities or the opportunity to work collaboratively with more experienced researchers in another country, HPC-EUROPA2 provides a valuable service to the HPC community which is complementary to PRACE. It is interesting to note that several of the current PRACE and DECI users benefited from transnational research visits under HPC-Europa and its predecessor projects, early in their research careers;
– Ignoring the absence of a European dimension of the TeraGrid, it is an interesting model that should continue to be investigated and analysed by PRACE in its further efforts to upgrade a European Tier-1 resource exchange model. At the moment, TeraGrid appears to be analogous to PRACE Tier-0, in that it operates a single, fully integrated Peer Review system with contributed resources, which do not, as yet have a comparable element of resource exchange to DEISA. TeraGrid also operates within a single country making it administratively more simple to organise and manage, although resources come from a number of different agencies and organisations, making its governance model more straightforward than that of DEISA or PRACE. However, it is likely that the TeraGrid model will evolve as TeraGrid itself is replaced by a follow-on project (XD) and so will, in parallel with PRACE, be considering changes to its resource exchange model. Continuing discussions with XD on resource exchange models would be mutually beneficial and could facilitate interchange between the two continents.

4.3 Recommendations

It is highly desirable for PRACE to continue the review of resource exchange models used by other e-Infrastructure initiatives and projects, particularly these models which are still evolving. This will provide valuable insights into alternative approaches with a view to improving or adapting in future the PRACE Tier-1 model to better meet the future needs of the European HPC ecosystem. Some interesting perspectives and essential replication leading to extension and additional strengths to the emerging PRACE resource exchange model can be obtained by:

– TeraGrid – discussion over the past few years about moving to a different exchange rate mechanism. For instance, at the moment an applications benchmark scheme is not

in global use, but the application benchmark numbers, if known, can be used if the application is moved between systems. Another example that is necessary to be considered is provision of resources owned by one organisation. It can be a "pooled resource" for EU funded projects. For instance the National Science Foundation (NSF) owns most of the resources provided by TeraGrid, with some also coming from other agencies. The science case for already awarded NSF projects is not re-assessed but only the appropriateness of the computational approach. In addition, TeraGrid allocations mechanism is organised around quarterly review requests and can span multiple years not only for 1 year maximum as it is applied in most of the other models. Following the discussion about inclusion of data storage in the exchange rate mechanism in future, also bandwidth and network.

- EGI – access granted to Virtual Research Communities, by way of Virtual Organisations (VOs), and these communities have to negotiate with resource providers, in general the National Grid Infrastructures (NGIs), for the allocation of resources.
- HPC-Europa2 – provision of enhanced training and consultation activities. In addition, investigation of the in-kind contribution mechanism (valuation and reporting) as around 20% of the project financed cycles is in-kind contribution. Last but not least, gaining some perspectives about submission application process from submitting an application through to the delivery of the final report which is based on a web-based management tool. Each researcher registers himself and completes the application form electronically. Both the technical evaluation and the host evaluation forms are completed on line.
- NDGF - building "pooled resources". Each participating country allocates a set of resources to a common pool from where the users (research groups) can select the services they wish to use. Resource exchange does not happen by direct trading between users, centres, or countries. The RT project identified the "unit of exchange" question as one of the most difficult issue to tackle when planning to build a resource exchange market, in particular when resources exchanged are of different types (e.g. comparing the value of CPU hours to expert work is not as straightforward as comparing the value of CPU hours on one machine to CPU hours on another machine). According to the conclusions of the project, the most feasible option seems to measure the value of each service in real currency (e.g. €) and define a conversion factor to produce credit units.
- HP-SEE - feedback is undertaken via mechanisms including project reports, surveys, documentation of experience in the wiki, list of publications, helpdesk system are some of the defined way for collecting feedback from the applications. In addition, in case of preparatory access applications are given resources to the so-called "home" cluster to test their suitability for porting and optimization to the HPC resources available in the region. The "home" cluster offers limited resources for the above purposes. After that each application sends their request for normal access to the ARC.
- LinkSCEEM – preparatory access deployed allowing researchers to apply for code scalability testing and also support for code development and optimisation as well as the creation of an HPC eco-system for the scientists in the region and enable them to integrate their small cluster resources with Teraflops computers (supporting vertical integration issues) coupled with the appropriate training and technical support.

Based on the existing models and practices, an evolutionary approach can be recommended combining PRACE members' readiness to provide HPC Tier-1 resources and user Tier-1 needs and expectations. The results of the review of existing models suggest that the DEISA

DECI model comes closest to meeting the immediate PRACE requirement to identify a starting point for a model for PRACE to use for the pilot Tier-1 call. The model will undergo further and more detailed evaluation during the call to assess its longer term suitability for PRACE, given the possible future synchronisation with Tier-0 process and increases in the scale and dimensions of subsequent Tier-1 calls. We will also continue on-going discussions with the architects of other resource exchange models, to collaborative on the joint development of solutions to the increasingly complex challenges of resource exchange, so that we can meet our ultimate goal of facilitating seamless access for researchers to a full range of facilities.

5. Input from User Questionnaires

During the past two years, DEISA has undertaken several user surveys, contacting the DECI users who participated in DECI-4 and DECI-5 calls. PRACE WP7 also undertook a survey of HPC users of the national facilities of the PRACE partners.

Output from these questionnaires was studied as it casts some light on the likely future demand for transnational HPC access and on user requirements in this area. This information is important in providing input about the future requirement of a Tier-1 exchange programme.

This analysis will be undertaken in two stages. Firstly, to distil user requirements then to assess whether these requirements should be catered for by PRACE, with recommendations on the role of Tier-1 resource exchange in meeting user requirements.

History of DECI

The DEISA Extreme Computing Initiative (DECI) is a scheme through which collaborative projects, headed by European scientists, can apply for single-project access to world-leading national (Tier-1) computational resources in the European HPC infrastructure for a period of up to 12 months per project.

DECI was introduced in 2005 to enable European computational scientists to obtain access to the most powerful national computing resources in Europe regardless of their country of origin or work, and to enhance DEISA's impact on European science and technology at the highest level. Through an annual call, a number of capability computing projects are selected by peer-review on the basis of innovation and scientific excellence. To date, six DECI calls have been held.

Successful projects are given access to the national HPC resources in the DEISA infrastructure (on HPC architecture selected for its suitability to run the project's codes efficiently) and are offered applications support to enable them to use it productively. The number of proposals received has grown, from 51 at DECI-1 to 122 at DECI-6, with particularly rapid growth over the past two years (during which access to PRACE Tier-0 systems via PRACE calls has also been available, suggesting that the availability of Tier-0 resources has stimulated rather than suppressed demand for Tier-1 access, especially as a step for accessing Tier-0 resources.) The number of CPU cycles requested per project has also grown steadily since DECI's inception.

DECI has been of key importance in building a European HPC user community, supported in their use of top-level HPC facilities by applications experts from leading European HPC centres. More than 330 different researchers from 25 European countries have participated in the scheme to date as Principal Investigators or as Co-Investigators, along with collaborators from four other continents. The scheme aims to enhance Europe's international standing in

science, by providing access to world-class resources which are complementary to resources available locally (i.e. at the researchers' national facility). DECI is a very successful instrument for the European Union research policy of enabling smaller countries to take part in research at the cutting-edge and enables all European researchers to participate in internationally competitive computational science.

DECI also aims to facilitate a better understanding of the likely requirements of future users of the Tier-0 systems by collecting real use-case information about what European computational scientists want and about the differences between usage of national and European resources and facilities.

The DECI-6 call attracted 122 proposals for challenging European computational science projects, requesting more than half a billion processor(-core)-hours and asking for significant application support. The call was very heavily oversubscribed both in requests for CPU and for applications enabling effort. 91 million processor (-core)-hours were available for distribution on 16 European HPC systems, and these were allocated to 55 projects, which have been given production access to the infrastructure between 1 October 2010 and 30 April 2011.

Figure 1 shows how both the median and the average amount of CPU requested by projects has grown year on year much faster than DECI provision can satisfy. While available cycles are increasing, so are European users' expectations and requirements. It is essential that we take steps to meet their burgeoning requirements if they are to remain inter-nationally/continentally competitive.

The growing divergence between the average and the median CPU requested reflects the increasing number of larger, collaborative projects applying to DEISA for computational and applications enabling resources. For example, particle and plasma physics projects (where there is a long tradition of cooperation at the European and international level) generally request many more resources on average per project than other application areas. The materials science community, however, is much more disparate and submitted a much larger number of smaller project proposals.

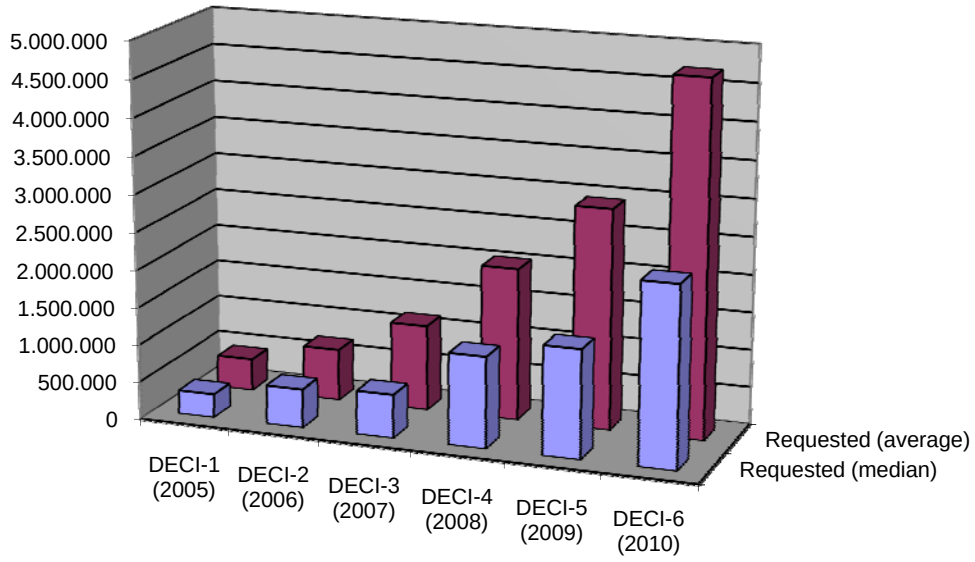

Figure 1 Average and median amount requested during the lifetime of DECI (note that all resources requested are expressed in normalised DEISA (processor-)core hours)

The following section gives an overview of the projects running on the DEISA infrastructure during 2010-2011 to show the range of science supported.

The 55 DECI projects selected for DECI-6 were self-categorised (by the PIs) into six broad scientific disciplines – Astronomical Sciences, Biological Sciences, Earth Sciences, Engineering, Materials Science and Plasma/Particle Physics. Overall, only 45% of the proposals were able to be accepted and the amount of resources awarded to many projects was scaled back to enable as many projects as possible to be supported.

There were 151 investigators named in the 55 accepted projects. Of these investigators, 106 (70%) came from countries with a DEISA partner site, 30 (20%) from other European countries and 15 (10%) from large outside of Europe (China, Israel, Japan, Saudi Arabia, USA). This suggests that successful project proposals were generally based on collaborative European science. All of the 58 PIs (three projects had two PIs) were Europeans, offering further proof that the research being undertaken focuses on European research priorities, with 49 (84%) of the PIs coming from countries with DEISA sites and 9 (16%) coming from other European countries.

In an effort to find out more about the degree of co-operation and collaboration within DECI projects, we analysed the information on investigators in further detail. The number of investigators in DECI-6 projects varied from one to seven. During the six years in which DECI has been operating, the average number of investigators has risen steadily from 1 to 3.

Finally, we looked at the number of countries collaborating in successful proposals. This showed that 38 (69%) of the projects involved scientists from more than one country, with 14 (25%) involving scientists from three or more countries. This indicates that DECI is

supporting pan-European scientific collaboration, and that the Tier-1 infrastructure is attractive to European researchers. Of the 364 scientific investigators who have used the DEISA infrastructure via DECI, 281 (77%) have been involved in one project to date with a further 83 (23%) involved in two or more projects. This shows that DEISA is being successful in reaching out to new groups and collaborations. It also suggests that DEISA resources are recognised as being complementary to national facilities, with applications only being made to DEISA when the circumstances warrant it.

Overall, the statistics which we have collected suggest that DECI is being successful in attracting high quality collaborative proposals involving scientists from more than one European country, and in involving partners from outside Europe where they add value to the collaboration.

The 55 DECI projects accepted specified 51 different named codes between them which they wanted to use plus two unnamed "own" codes. Of the 50 codes which specified their communications method, 39 (78%) of the codes were parallelised with MPI, with a further 8 (16%) having hybrid OpenMP/MPI implementations and three (6%) using MPI and pthreads. Of the 51 named codes, 45 were used by just one DECI project, with a further three being used by two projects, two (Chroma and VASP) by three projects and one (NAMD) by five projects. In particular, the large number of different codes used by DECI projects indicates the wide variety of science being undertaken at the EuropeanTier-1 level and reinforces the need for close partnership with users to help them to achieve good performance via applications enabling and code tuning. There is obviously a strong demand for the sort of complementary computing support offered by DECI which crosses scientific domains.

It will be also interesting to analyse if some DECI projects migrated to Tier-0 and were successful in PRACE AISBL calls for Tier-0.

Input from DECI questionnaires

DEISA designed a questionnaire which was sent to all the researchers who used the DEISA infrastructure in 2009. The questionnaire had two main purposes, firstly, to provide more feedback on the quality of the DEISA user experience (including more detailed information about how the services provided were perceived and used) and secondly to help DEISA to articulate the ways in which the project added value to computational science research in Europe. The questionnaire remained open until 28 March 2010. The number of responses received was 105, with responses from users of all of the supercomputers in the DEISA infrastructure. There was a good mix of responses covering the whole spectrum of DEISA users – 40 respondents were PI or CoIs (38%), a further 40 were researchers (38%) and the remaining 25 (24%) were research students. The vast majority (80%) were first time users of the DEISA infrastructure with 20% having used it on a previous project at some point during the preceding four years.

The activities where DEISA appeared to have had the greatest impact were in "Scaling codes to run on larger HPC systems" (44% rating DEISA as having had a high impact) and in "Obtaining access to expert technical or application enabling support" (23% rating DEISA as having had a high impact).

However when we look at both high and average impact, we see that 72% of respondents reported that "Experimenting with new technology and/or alternative architectures" had had a high or average impact for them.

These statistics suggest that the primary benefit which DECI delivers is to enable researchers to scale codes to run on larger HPC systems. The applications enabling support which is

provided is an important and highly regarded feature of the service while the opportunity to experiment with new architectures is also highly valued.

DEISA also asked respondents what impact the use of the DEISA infrastructure and involvement with DEISA had on their research. The table below shows how researchers felt that DEISA had benefitted their research.

DEISA has enabled me to….	Percentage replying "Yes"
Work on research problems that I could not address before	70%
Accomplish research tasks more quickly	61%
Do more accurate, higher quality research	54%
Access resources for my research faster or better	53%
Produce more research output per year	51%
Produce, process or analyse data faster or better	46%

Table 2 Impact of use of the DEISA infrastructure.

From table 2, we see that more than half of the respondents reported that DEISA had enabled them to work on research problems that they could not address before, to accomplish research tasks more quickly, to do more accurate, higher quality research and to access resources for their research faster or better.

Overall, 97% of respondents said that the availability of DEISA or similar resources for their research work was very important or important. For a majority of respondents (53%) it was even "very important". Only one respondent felt that the lack of DEISA or similar resources would not impair their research at all, with 9% saying that it would impair their research "totally", 70% saying that it would impair their research "much" and a further 20% saying would impair it "a little".

These results show how important and valuable DEISA has now become to the European computational science research community.

Input from the second DECI questionnaire

A second questionnaire (formulated with input from this PRACE workpackage) was sent to all the DECI researchers who used the DEISA infrastructure in 2010. The questionnaire had two main purposes, firstly, to provide DEISA with more feedback on the quality of the DEISA user experience (including more detailed information about how the services provided were perceived and used) and secondly to help to articulate the ways in which DEISA adds value to computational science research in Europe (with a view to formulating a plan for continuing DECI in PRACE).

The questionnaire remained open until 21 January 2011. The number of responses received was 55, with responses from users of all of the supercomputers in the DEISA infrastructure and spanning all of the major scientific disciplines. We had a good mix of responses covering the whole spectrum of DEISA users – 30 respondents were PIs or Co-Is (55%), a further 15 were researchers (27%) and the remaining 10 (18%) were research students. The majority (62%) were first time users of the DEISA infrastructure.

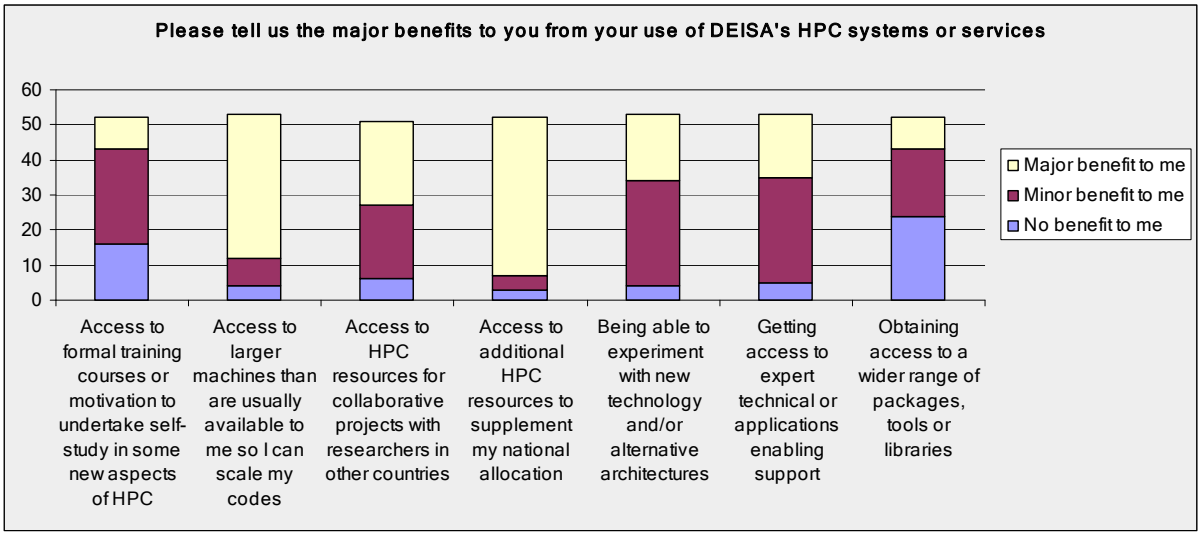

Figure 2 Major benefits of using DEISA's HPC systems from the point of view of users.

As can be seen from figure 2, the activities where DEISA appeared to have had the greatest impact were in "Access to larger machines than are usually available to me so I can scale my codes" (77% rating DEISA as providing a major benefit) and in "Access to additional HPC resources to supplement my national allocation" (87% rating DEISA as having had a major benefit).

However when we consider both major and minor benefits, we see that 92% of respondents reported that "Being able to experiment with new technology and/or alternative architectures" was of benefit4, 91% of respondents reported that "Getting access to expert technical or applications enabling support" was of benefit and 90% reported that "Access to HPC resources for collaborative projects with researchers in other countries" was of benefit.

These statistics suggest that the primary benefit which DECI delivers for "newcomers" is to enable researchers to scale codes to run on larger HPC systems. The applications enabling support which is provided is an important and highly regarded feature of the service. The responses also clearly indicate that, in the vast majority of cases, DEISA resources are complementary to national resources and that there is strong interest in experimenting with new technologies and alternative architectures. The role which DECI plays in making resources available to collaborative projects is also appreciated.

We also asked respondents what impact the use of the DEISA infrastructure and involvement with DEISA had on their research. The table below shows how researchers felt that DEISA had benefitted their research.

[4] In this context, "new technology" often means "technology that is new to me",

DEISA has enabled me to….	Percentage replying "Yes"
Work on research problems that I could not address before	76%
Accomplish research tasks more quickly	74%
Do more accurate, higher quality research	66%
Produce more research output per year	57%
Produce, process or analyse data faster or better	49%
Access resources for my research faster or better	45%
Have more publications or conference proceedings accepted	25%

Table 3 Impact of use of the DEISA infrastructure.

From this, we see that more than three-quarters of the respondents reported that DECI had enabled them to work on research problems that they could not address before, and almost three-quarters stated that DECI had enabled them to accomplish research tasks more quickly. Almost two-thirds stated that DECI had enabled them to do more accurate, higher quality research and over half stated that DECI had enabled them to produce more research output per year.

One quarter stated that more publications or conference proceedings had been accepted but this is most likely to be due to the inevitable time-lag between producing results and publishing them, and this percentage may increase due to the fact that using DEISA allowed researchers to make more accurate and higher quality research scored 66%.

All respondents stated that DEISA resources had been either very important (64%) or important (36% to them). All respondents stated that the lack of a DEISA (Tier-1) infrastructure would impair their research to some degree, with 85% stating this would much or totally impair their research.

These results show how important and valuable DEISA has now become to the European computational science research community and fully justify the decision to continue DECI in PRACE AISBL. DEISA has also allowed scientists to work in larger computer systems and has also supported code development for achieving larger scalability, which is a very important factor to help scientists to make the transition from Tier-1 to Tier-0 computer systems and as such supports the integration of DEISA into PRACE AISBL.

Input from PRACE questionnaires

The PRACE project undertook a survey of Applications and Users Requirements for Tier-0 systems which closed in January 2011. It aimed at providing a snapshot of the current usage of HPC systems operated by PRACE partners and at collecting information about usage patterns, behaviour and requirements of current and potential users of PRACE Tier-0 systems. The survey sample was self-selecting in that PRACE partners contacted any users who might potentially have an interest in using PRACE Tier-0 systems to invite them to take part. The peak performance of the 27 HPC systems used by respondents varied from 1.02 Pflop/s peak to 6.3 Tflop/s peak with a median of 59 Tflop/s peak. For comparison purposes, the current DEISA Tier-1 systems are all above the median, so the survey can arguably be said to tell us as much about future demand for larger Tier-1 systems as about demand for Tier-0 systems.

The survey results indicated that although job sizes in terms of number of cores have increased since the last PRACE survey in 2008, this increase has not kept pace with the number of cores in the machines so that machines are now typically utilised by a larger number of concurrent jobs. There was a strong correlation between job size and system size,

with over half of the total CPU being used in jobs of more than 8192 processors on a small number of large HPC systems. In order to estimate how many of these national users were potential users of PRACE Tier-1 or Tier-0 systems, we looked at the results of the questions about current and desired/required scalability (see table 4).

No of cores	Actual scalability (% of respondents)	Desired or required scalability (% of respondents)
Not yet parallel	5%	1%
2-128	38%	20%
129-512	21%	20%
513-2048	17%	19%
2049-8192	7%	13%
> 8193	12%	27%

Table 4 Current and desired scalability survey

The results show a clear desire from most users to increase the scalability of the applications which they are using, but in most cases the desired or expected scalability is relatively modest. These findings seem to indicate the following:

- Access to PRACE Tier-1 systems (and applications enabling assistance to scale codes by an order of magnitude or more) would be welcomed by a large segment of Tier-1 users and in some cases even of the potential Tier-0 user community;
- Tier-1 has an important role to play in providing a ramp to Tier-0. The majority of users are thinking in terms of incremental performance improvements. However the proportion of users whose codes already scale beyond 8192 processors is quite small (12%), although there is a similar sized pool of users (an additional 14% of the European HPC Tier-1 user base) for whom this is currently the required or desired scalability. If PRACE were able to target more of these users (who are probably mostly in the 512 processors to 8191 processors categories) and provide them with appropriate support and Tier-1 access, then DECI could become a very effective instrument for providing a ramp to Tier-0;
- Improving scalability of codes will, of course, increase demand for both Tier-0 and Tier-1 resources. The survey identified that more than half of the total resources were used by the 12% of the users whose codes scale to more than 8192 processors. In the short term, targeting those users who require scalability to 8192 or more processors, using Tier-1 as a ramp, is likely to place further demands on the already stretched Tier-1 resources.

About a third of respondents, (36%) indicated that the time allocated to their research group was used by collaborators based in another country. Although it was not clear if these responses related just to resources allocated on their national HPC resources or also included resources from European infrastructure projects such as DEISA, HPC-Europa or PRACE, the replies confirm the DECI experience that computational science research is becoming increasingly collaborative.

Finally, 63% of respondents indicated that the availability of intermediary machines (between their existing production facilities and Tier-0 systems) would make it more likely that they would, in time, apply for Tier-0 access.

Summary and conclusions of questionnaire results

The questionnaires referred to in the sections above were obviously not designed to provide user input on the question of what a model for European HPC resource exchange should look like. Nor has the European community of HPC users been surveyed in a systematic and consistent way. However, it does seem likely that they provide a useful insight into the current and future role which European Tier-1 access and services play in the HPC ecosystem and allow us to draw some conclusions about the likely future demand for Tier-1 systems and services and the ways in which European resources are likely to be used. The surveys also show that access to Tier-1 systems is important for some users as an intermediate step for preparing future access to Tier-0 systems.

6. PRACE Distributed European Computing Initiative (DECI)

The future Tier-1 exchange programme is an important strategic development for PRACE AISBL. When WP4.3 was originally proposed, this deliverable was intended primarily to provide information on which to make informed decisions about a future, but as yet undefined, Tier-1 resource exchange programme.

However PRACE-2IP has the stated objective of integration of Tier-1 resources, services and user communities within the PRACE RI with a specific focus on re-using the successful components of the three DEISA projects. DECI calls (DEISA Extreme Computing Initiative), the current single-project transnational access scheme for Tier-1, will be continued with the same acronym but now meaning Distributed European Computing Initiative.

As a result, it was agreed in PRACE 1IP-WP4 that task 4.3 could initiate a first DECI pilot to be implemented under PRACE AISBL, as part of its work on the development of a Tier-1 exchange model.

6.1 Initiating a DECI pilot under PRACE AISBL

The purpose of this pilot is to implement an initial DECI call in PRACE AISBL, based largely on the existing DEISA practices and procedures with some modifications prompted by our findings to date. Once the pilot is completed, the feedback and collective experience can be used to refine the resource exchange model for future calls.

After seeking feedback and guidance from the PRACE BoD and PRACE 1IP/2IP Management Board, the following points were agreed:

1) In order to ensure the continuity of the DECI calls to the scientists and to avoid gaps in the bidding process, the first PRACE-DECI pilot should be opened in May 2011, and implemented in November 2011, coinciding with the third Tier-0 regular call.
2) Resources for the DECI call should come from PRACE-2IP which will be formally responsible for implementing the call. Support for the first Tier-1 call under PRACE AISBL, in particular during the early phase when applicants are writing their proposal (advice on what machines to ask for, answering queries, etc.) should also be planned by PRACE-2IP.
3) Given the time constraint, it was suggested that the pilot could primarily rely on the partners who were presently involved in DEISA 2. These partners have a long experience of working together, which will ensure a fast and smooth transition to the

PRACE environment. Additional partners were nevertheless encouraged to join the pilot.

4) The processes and practices related to the allocation of resources (5% rule) and peer-review (national level) will remain the same as in DEISA for the first pilot.

6.2 Description of the DECI pilot

The new DECI call was announced on 13-14 April 2011 at the DEISA-PRACE Symposium in Helsinki, Finland, and advertised through different channels after this event. The paragraphs below summarise the main features of the call.

Scope of the call

The scope of the new DECI remains very similar to the DEISA Initiative: DECI enables European researchers to obtain access to the most powerful national (Tier-1) computing resources in Europe, regardless of their country of origin or work and to enhance the impact of European science and technology at the highest level.

Proposals must deal with complex, demanding, innovative simulations that would not be possible without Tier-1 access. In addition to offering access to computing resources, applications-enabling assistance from experts at the leading European HPC centres is offered to enable projects to be run on the most appropriate Tier-1 platforms in PRACE.

Projects supported by DECI will be chosen on the basis of innovation potential, scientific excellence and relevance criteria. Priority will be given to proposals that promote collaborative research, either at a cross-national or cross-disciplinary level.

Tier-1 resources

It should be noted that, at the time of writing this deliverable, a larger number of systems are offered in this new DECI call compared to the last call. For Tier-1, applicants should indicate a preference for a class of architecture. Four classes have been defined below.

- Cray XT4/5/6 and Cray XE6 – three large Cray XE and XT systems are available at EPCC (UK), SNIC-KTH (Sweden) and CSC (Finland). The largest of these machines has a peak performance of 360 T/flops and a total of 44,544 cores.
- IBM Blue Gene/P – three BG/P systems are available at IDRIS (France), RZG (Germany) and NCSA (Bulgaria). The largest of these machines has a peak performance of 139 T/flops and a total of 40,960 cores.
- IBM Power 6 – three IBM Power 6 systems are available at RZG (Germany), SARA/NCF (The Netherlands) and CINECA (Italy). The largest of these machines has a peak performance of 98 T/flops.
- Clusters – PRACE offers access to eight clusters at FZJ (Germany, Bull Nehalem cluster), LRZ (Germany, Xeon cluster), HLRS (Germany, NEC Nehalem cluster plus GP/GPU cluster), CINES (France, SGI ICE 8200), BSC (Spain, IBM PowerPC), CINECA (Italy, Westmere plus GP/GPU cluster), PSNC (Poland, Bullx plus GP/GPU cluster) and ICHEC (Ireland, SGI ICE 8200). The largest cluster has a peak performance of 267 T/flops and a total of 23,040 cores

Timeline

The new DECI call will open on 2nd May 2011 and close on 22nd June 2011. Access will be awarded for a period of 12 months, beginning 1 November 2011. This Tier-1 call will have the same timeline as the Third Tier-0 Regular Call.

7. Adapting the DECI calls to PRACE AISBL context

One of the purposes of the pilot is to provide us with additional information on which to make recommendations relating to the future PRACE AISBL Tier-1 exchange programme. After the 1st DECI pilot has been initiated, WP4.3 will address a number of issues that are important for the integration of the Tier-1 exchange programme in PRACE AISBL. These issues will be addressed in detail in the second deliverable (D4.3.2, M24) using the experience from the pilot to inform the decisions.

he migration of DECI into the PRACE AISBL environment will require further modifications of the programme to be consistent with the continuing evolution of DECI's role in the European HPC ecosystem. This includes consideration of the relationship between Tier-0 and Tier-1 (now that the national resources provided at the Tier-1 level no longer represent the apex of the pyramid) and the expansion of Tier-1 resource provision to include a wider number of national partners, a greater range of architectures and a larger discrepancy between the biggest and smallest Tier-1 facilities.

For the benefit of all stakeholders, including potential users of the PRACE AISBL resources, the national funding bodies contributing Tier-1 resources and the Tier-1 centres operating the HPC facilities, it is important that the PRACE Tier-1 resource exchange procedures adhere to a number of basic principles. This means planning and conducting the scope of the resource exchange activities using proper methodologies to promote transparency, accountability and objectivity, by adoption of a suitable framework to this end.

Tier-0 cycles providing model brought up the issue of VAT and taxation for the AISBL and the Hosting Members for its potential of being perceived as value-added service to customer by the Belgian and local tax authorities. The extensive analysis of the legal implications, both at AISBL and national government level, led to the indication that no VAT taxation would arise if:

- There isn't any transfer of goods and assets between Hosting Members and PRACE
- The Hosting Members only transfer the decision power on the resource assignment to PRACE AISBL

Whether the Tier-1 resource exchange model will comply with the above conditions is not yet clear. Moreover, the number of national sites that will contribute to the resource exchange will be certainly larger than Tier-0 and more dynamic.

Another point that could be addressed in the future is to see whether regional clusters, such as represented by HP-SEE or LinkSCEEM, can play a role or have an impact on the exchange model (e.g. jointly providing/operating Tier-1 resources). This will be explored with the resources providers from these regions.

It is important that Tier-1 resource exchange model takes into account the VAT taxation issues from its early stage of development taking advantage of the legal support provided at project level.

A few issues that will need to be considered have been already identified. They relate to the synchronisation of Tier-1 and Tier-0 calls, the peer-review process, the resources allocated by partners, and the number of partners involved in the exchange programme.

Synchronisation of Tier-0 and Tier-1 calls

This year, the timeline of Tier-1 and Tier-0 calls will be for the first time synchronised, providing users with the possibility to apply to whichever call they find as the most suitable for their project, although the application tool will remain in practice separated.

Opening the two calls at the same time is already a big step towards greater visibility for the users regarding the types of Tier-0 and Tier-1 resources that are available through PRACE. Joint dissemination activities (e.g. press release and other advertising) will also contribute to this goal.

Ultimately, the PRACE AISBL may consider further integration between the two calls, by merging e.g. the application and the evaluation processes. In order to achieve this, further investigation of the benefits of this integration should be conducted as well as examining various mechanisms to handle the applications and to direct users to the most suitable resources.

Here is a brief list of a few points to consider:

- What would be the benefits of a better synchronisation of the two calls?
- What should it cover (proposal submission, peer-review process, etc.)?
- What are the advantages and disadvantages of a single application process which enables resources to be requested and awarded on multiple resources, across different levels of the HPC ecosystem?
- How could a resource exchange model facilitate this?

Peer-review process

Peer-review processes are today different for accessing Tier-1 and Tier-0 resources. In the context of the integration of Tier-1 resources into PRACE AISBL, the following questions will need to be considered:

- Given that Tier-1 resources are contributed nationally, how should peer review of DECI proposals be handled?
- Should all DECI proposals be peer-reviewed by a single body or should we continue with a distributed model of national peer review which will become increasing difficult to administer transparently and objectively as the number of partners increases?
- If the projects are reviewed nationally, could we accept projects from national calls which could be placed into the DECI resource pool for exchange?
- How should projects with multiple partners from different countries be peer-reviewed?
- How can the issue of double evaluation be addressed for projects requesting resources at both Tier-0 and Tier-1?
- How can we ensure transparence and fairness to the users with different peer-review processes?

Resources allocated by partners

In the light of the expected demand for Tier-1 resources (based on DEISA's evidence of demand rising faster than supply) are there any resource exchange strategies which could motivate national funding bodies to increase their resource contributions to PRACE AISBL Tier-1 calls? Could or should we evaluate and include additional contributions such as data storage, visualisation facilities, software licences etc. into the resource exchange model? What is the role of a resource bank in facilitating partners opting in or opting out of particular calls based on their circumstances?

Optimum number of partners to be involved in each call

Given the large number of partners in PRACE AISBL, is the preferred model for all partners to be involved in every call or should we expect partners to contribute a larger faction of resources to a smaller number of calls (this seems to be more important for partners with reduced amount of resources)? Will/can there be different roles in the programme taking into consideration the heterogeneity of the partners (in terms of size, capacity, type of architecture, etc.)?

8. Conclusion

The development of a cross-national model of resource exchange for access to Tier-1 resources is an ambitious undertaking which is nevertheless essential if Europe is to develop seamless support for computational science at all levels of the HPC Ecosystem. In the HPC arena, nothing of this scale and complexity has been attempted to date. There are therefore few examples of good practice for PRACE to follow. Perhaps not surprisingly, given the large number of countries in Europe and the structuring effect of EC funding, the projects with models which seem to offer the most in the way of guidance to PRACE are European – DEISA and HPC-Europa.

The role of this task in PRACE-1IP is to develop and pilot practices to stimulate Pan-European HPC access. In support of this objective, a decision was taken to proceed with a pilot call for DECI in PRACE which opened in May 2011 (some five months before the start of the PRACE-2IP project in which Tier-1 services will be integrated into PRACE).

The early scheduling of this call says much about the importance which PRACE attaches to the integration of Tier-1 into the PRACE infrastructure. Equally, it is widely understood that the selection of an appropriate resource exchange model which meets the needs of the different stakeholders will be key to its long term sustainability and to its growth.

This deliverable has documented the work which has been undertaken to date in setting the scope for the Tier-1 resource exchange model through a definition of both a Tier-1 centre and a Tier-1 system. It has also provided an overview of the research into existing models of resource exchange and into user requirements via the results of a number of user questionnaires.

Over the course of the next few months, PRACE will run a pilot, the main purpose of which is to assess how best to integrate DECI (and the DEISA Tier-1 services) into the PRACE infrastructure. Initially, the pilot will be modelled for simplicity's sake, on the DEISA DECI resource exchange model. It will, however, provide this workpackage with valuable feedback on some operational aspects of the DEISA model in its new environment while we continue to collect information and to refine our requirements.

One of the key questions for PRACE is to determine the medium to long term role of Tier-1 within the infrastructure as this will have a bearing on the optimal design of the resource exchange model.

We believe that this document will play an important part in informing the PRACE AISBL. It will be followed up by a second deliverable, produced in M24 which will document the PRACE resource exchange model which has been adopted. Between now and that date, there is a large amount of work to be done in further defining and agreeing the framework for evaluating resource exchange models (based on the AISBL's more detailed requirements), on analysing the medium to long term demand for both Tier-0 and Tier-1 resources (to clarify the future role of Tier-1) and to examine the potential of a resource bank to manage supply and demand more effectively.

Annex 1

Template and Notes from the meetings with Person of existing resource exchange initiative or project

1. Template

Notes from meeting (or email exchange) with Person or Persons contacted, of Project Name on date(s)

A sentence about who the person or persons contacted are and their roles on the project.

What is Project Name?

A short description of the project and its objectives.

For the following questions, if a resource exchange system is not yet in operation, please collect information about intentions e.g. What resources will be exchanged/ instead of What resources are being exchanged?

What resources are being exchanged?

Where do the resources being exchanged come from?

What is the unit of exchange?

Does Project Name have a resource bank?

Does Project Name use incentives to encourage use of particular resources?

How are resources allocated?

How is feedback collected (about allocations) from users?

Does Project Name offer preliminary allocations?

Prospects for further collaboration with PRACE?

2. DEISA

Notes from email exchange with Hermann Lederer, of DEISA on 15 February 2011

Dr. Hermann Lederer is the Deputy Director of RZG, the co-ordinating partner in the DEISA and DEISA2 projects. He has had extensive involvement in the definition of the DEISA resource exchange model.

What is DEISA?

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

DEISA deploys and operates a persistent, production quality, distributed supercomputing environment with continental scope. By extending the European collaborative environment in the area of supercomputing, DEISA has paved the way towards the deployment and operation of a persistent cooperative European HPC ecosystem, as suggested by ESFRI, and provided a turnkey operational solution for it.

DEISA2, funded by the European Commission in FP7, continues to develop and support the pan-European distributed high performance computing infrastructure established since 2002 within the predecessor project, DEISA1 that was funded in FP6. The DEISA infrastructure is based on the integration of thirteen national supercomputing centres from nine European countries. The dedicated DEISA network has been realized through GÉANT2 and the NRENs.

What resources are being exchanged?

The project is exchanging CPU cycles between thirteen national supercomputer centres so that researchers from any European country can obtain peer-reviewed access to the most suitable architectures and machine configurations for their science. The Applications Task Force provides high level applications enabling, and operations and user support is given by respective European teams, staffed by all of the partners. Data storage is also provided to users for the duration of their projects, although this is not formally accounted for in the resource exchange process.

Where do the resources being exchanged come from?

Each participating site contributes a minimum of 5% of the cycles from its national or large regional HPC systems, by agreement with its funding body or research councils. Almost all of the participating machines are in the Top 100 worldwide.

What is the unit of exchange?

The unit of exchange is the DEISA processor (core) hour. Local CPU hours are converted to DEISA processor hours and all allocations are made in DEISA units. The conversation rates are set in advance of allocations being made. The DEISA benchmark suite is run on each new system or upgraded system to provide a measure of applications performance. Geometric means are used (discarding the two highest and lowest results) to come up with a measure of performance relative to the other machines in the DEISA infrastructure. For simplicity, machines of the same architecture and similar configurations are usually allocated the same conversion factor.

Does DEISA2 have a resource bank?

Resources are contributed by partners for a 12 month period. Within that period, resources are allocated to DECI calls and to virtual communities, usually in fixed proportions. There is not a resource bank per se, but sites can vary their contributions to different initiatives e.g.

concentrating on DECI or on Virtual communities or can bank additional time to provide additional resource for projects of national importance.

Does DEISA use incentives to encourage use of particular resources?

No. However, individual partners sometimes offer incentives to DEISA users who have been allocated to run at their site (identical to incentives offered to national users) to reward usage of a particular queue or machine partition etc.

How are resources allocated?

Proposals undergo both a technical review and a scientific peer-review. Based on these reviews and their rankings, taking into account user expressed preferences wherever possible, projects are allocated to machines by a committee. DEISA is not always able to offer users the full allocation which they have requested. In cases where users' allocations are scaled back, this is done with reference to both the scientific and technical evaluations.

How is feedback collected (about allocations) from users?

DEISA has an Applications Task Force (for DECI projects which are up to 12 months in duration) and Virtual Community Ambassadors (for virtual communities who are given access for a longer period of time by agreement). Each project therefore has a named person who liaises with them on all aspects of their project, including allocations.

Does DEISA offer preliminary allocations?

No, DEISA does not offer preliminary allocations. However, each of the DEISA machines is benchmarked against a suite of applications codes, to provide detailed information about the performance of each system with the same codes so we have existing information about the performance of different classes of codes on different architectures.

Scientists who submit proposals almost always have access to the national HPC systems of their own country.

In the unlikely event that a project is assigned to a resource which the PI finds is not optimal, DEISA has two possible routes for redress. The project can be wholly or partly moved to another machine in the DEISA network, usually via a swapping of cycles with another project. Alternatively, the project can be awarded additional resources to compensate for the perceived deficiency.

Persistence through PRACE?

The integration of Tier 1 and Tier 0 into a common European HPC ecosystem is a major feature of PRACE 2IP. DEISA services and DECI will be incorporated into PRACE.

3. TeraGrid

Notes from meeting with John Towns, TeraGrid on 17 November 2010

John is the chair of the TeraGrid Forum (TGF) and Senior Associate Director at the National Center for Supercomputing Applications (NCSA)/University of Illinois.

What is TeraGrid?

TeraGrid is an open scientific discovery infrastructure combining leadership class resources at eleven partner sites to create an integrated, persistent computational resource. It is therefore similar in concept and objectives to DEISA.

Using high-performance network connections, TeraGrid integrates high-performance computers, data resources and tools, and high-end experimental facilities around the United States of America.

What resources are being exchanged?

Currently, TeraGrid resources include more than 2 petaflops of computing capability and more than 50 petabytes of online and archival data storage, with rapid access and retrieval over high-performance networks. Researchers can also access more than 100 discipline-specific databases. With this combination of resources, the TeraGrid is the world's largest, most comprehensive distributed cyberinfrastructure for open scientific research.

Where do the resources being exchanged come from?

TeraGrid is coordinated through the Grid Infrastructure Group (GIG) at the University of Chicago, working in partnership with the Resource Provider sites: Indiana University, the Louisiana Optical Network Initiative, National Center for Supercomputing Applications, the National Institute for Computational Sciences, Oak Ridge National Laboratory, Pittsburgh Supercomputing Center, Purdue University, San Diego Supercomputer Center, Texas Advanced Computing Center, and University of Chicago/Argonne National Laboratory, and the National Center for Atmospheric Research.

Most of the resources provided are owned by the National Science Foundation, with some also coming from other agencies. Some of the hardware/resources are owned by third parties e.g. supplied by individual state; most is financed and owned by NSF.

What is the unit of exchange?

TeraGrid uses a normalised unit (Service Unit). This is based on a Cray XMP core-hour. Resources used are reported in local units which are then converted to normalised units[5]. Billing rates can potentially be different. Linpack is used to determine the exchange rate between machines. The advantage is that it is easy to determine the Linpack rating (from Top 500) and that the sites have historical date going back 20 years. However, there has been discussion over the past few years about moving to a different exchange rate mechanism. At the moment an applications benchmark scheme is not in global use, but the application benchmark numbers, if known, can be used if the application is moved between systems.

[5] SU Conversion Calculator - Use this form to convert service units (SUs) on one platform to the equivalent amount on another platform using the accepted conversion factors. This calculator uses a weighting factor based on the performance of resources against the HPL benchmark, which is used to determine the Top 500 list of HPC systems.

At the moment, data storage is only recorded and reported (not accounted for) but there have been discussions about including it in the exchange rate mechanism in future, also bandwidth and network.

Does TeraGrid have a resource bank?

TeraGrid does not have a resource bank. Allocations can be used at any time over the allocation period. There is a general assumption that there will be a flat burn of TeraGrid resources, but in practice, usage is in bursts.

Projects can request an adjustment to their allocation during the course of the project. I.e. they can request a move from one system to another.

Does TeraGrid use incentives to encourage use of particular resources?

TeraGrid does not offer particular incentives for users to use less popular machines. Some users are keen to use new resources, other prefer to stay on established machines.

Users will be allocated to a particular machine, which may or may not be their preferred choice.

How are resources allocated?

TeraGrid has one allocation process. Allocations can span multiple years. The review includes a technical review for feasibility and a scientific assessment. For applications that have received an NSF grant or equivalent financial support to fund the work-force from a national agency or foundation, the science case is not re-assessed but only the appropriateness of the computational approach. It is assumed that financial support for the scientific personnel to undertake the work is in place.

TRAC will also consider the progress made during prior allocations, including the publication of peer-reviewed papers and other communications within the community.

How is feedback collected (about allocations) from users?

TeraGrid undertakes regular user surveys, runs an annual meeting etc to get feedback from users.

Does TeraGrid have preliminary allocations?

Yes. Some resources are set aside for preliminary allocations. This is to enable projects to test their suitability for TeraGrid. This access is not peer-reviewed but applications are reviewed by an allocations committee. Only small amounts of access are available by this route to PIs who are new to TeraGrid.

Prospects for further collaboration?

Researchers of any nationality can apply to use TeraGrid. PIs must be a researcher or educator at a US academic or non-profit research institution.

TeraGrid would welcome the opportunity to discuss a mechanism for resource interchange with PRACE for future calls which would allow US-based researchers to access European facilities and European-based PIs to access US-based facilities. This would possibly be done via linked or joint calls. TeraGrid would also be interested in exploring standards based data sharing between Europe and the US.

The intention is to have a joint call between PRACE and TeraGrid in 2011 with shared resources.

4. EGI

Notes from meeting with Maurice Bouwhuis of SARA on 3 May 2011

Maurice Bouwhuis is Manager Relations and Innovation at SARA and is involved in the Dutch project BiGGrid, the Dutch NGI partner of EGI.

EGI - European Grid Initiative

EGI.eu is a foundation established under Dutch law to create and maintain a pan-European Grid Infrastructure (EGI) in collaboration with National Grid Initiatives (NGIs) and European International Research Organisations (EIROs), to guarantee the long-term availability of a generic e-infrastructure for all European research communities and their international collaborators.

Its mission is to enable access to computing resources for European researchers from all fields of science, from High Energy Physics to Humanities.

EGI provides access to resources, i.e. logical and/or distributed entities to be shared by end-users – for example CPUs, data storage, instruments and digital libraries. Resources are contributed by Resource Centres, also known colloquially as sites. The Resource Centre is the smallest localised or geographically distributed administration domain, where EGI resources are managed and operated.

What resources are being exchanged?

Access is granted to Virtual Research Communities, by way of Virtual Organisations (VOs), and these communities have to negotiate with the resource providers, in general the National Grid Infrastructures (NGIs), for the allocation of resources. Within EGI there is no model for the exchange of resources among resource providers.

Where do the resources being exchanged come from?
Not relevant.

What is the unit of exchange?
For accounting CPU usage is specified in kSpecInt and storage in TBytes. However, this is not used for exchanging allocations.

Does EGI have a resource bank?
No.

Does EGI use incentives to encourage use of particular resources?
EGI provides support to Virtual Communities for enabling applications on resources that use the Grid Middleware supported by EGI.

How are resources allocated?
In general resources are allocated through negotiations between Research Communities and resource providers. This can be different for NGIs, so there is no general rule on how to apply for access. Manpower resources for enabling effort are allocated centrally. These resources are provided by both EGI and NGIs on behalf of EGI.

How is feedback collected (about allocations) from users?

User surveys are used to gather user requirements.

Does EGI offer preliminary allocations?
Preliminary or preparatory access depends on the facilities offered by individual resource providers.

Prospects for further collaboration with PRACE?
There should be prospects for collaboration. There has been one meeting in 2010 to discuss the collaboration on operations.

5. HPC-Europa 2

Notes from Giovanni Erbacci and discussion with Sanzio Bassini (CINECA) of HPC-Europa 2 on April 28, 2011

CINECA is the co-ordinating partner of the HPC-Europa 2 project. Sanzio Bassini, HPC Director in CINECA, is the coordinator of HPC-Europa and Giovanni Erbacci is responsible for the allocation of the resources for HPC-Europa in CINECA.

What is HPC-Europa 2?

Under FP6, six main European HPC-centres associated together in the EU project HPC-Europa to provide integrated Transnational HPC-Access. The project offered the best way to integrate the HPC facilities and the expertises of these partners, effectively deploying a real virtual global infrastructure available for the benefit of users.

HPC-Europa 2 (www.hpc-europa.eu), funded by the EC in FP7, is a pan-European research infrastructure on HPC, that continues the direction established by HPC-Europa, with the objective to further advance the provision of integrated HPC computational services at the European scientific community working at the forefront of science. The HPC-Europa 2 core activity sees seven national HPC Centres involved in providing HPC Transnational Access (TA): CINECA (Italy), BSC (Spain), EPCC (UK), HLRS (Germany), SARA (The Netherland), CINES (France) and CSC (Finland). During the four years of activity of the project (January 2009 - December 2012) the transnational activity will provide HPC resources and services, specialist support, computational tutoring and scientific collaboration to more than 1,000 European researchers, greatly contributing to create an European HPC eco-system at the service of computational sciences.

What resources are being exchanged?

The project grants CPU cycles available at the seven national supercomputer centres involved in the TA activity.

The TA is granted to researchers from any European country on the basis of a peer-review selection of the applications.

Each TA centre is responsible for providing users with access to HPC facilities at their centre together with local scientific and HPC support. This service is enhanced and supplemented with a co-ordinated suite of training and consultation activities;

The Centres involved in the TA activity have extensive experience of working together to develop a highly effective co-ordinated approach for the management of the TA activities. Under this model the various TA activities are tightly integrated across the consortium, with a common application and selection procedure, a single selection panel and co-ordinated strategies for marketing, outreach and dissemination.

Where do the resources being exchanged come from?

The CPU cycles are provided by the seven participating HPC centres and the CPU cycles are financed by the EU through the HPC-Europa 2 Project. Furthermore the sites provide in kind a further amount of cycles from their national HPC systems around the 20% of the project financed cycles are in kind.

Furthermore the EU reimburses the researchers which succeed with a granted application, for travel and living expenses (until to a given threshold) sustained for staying at the hosting centre to use the HPC resources to realise the funded project.

What is the unit of exchange?

The unit of exchange is the Allocation Unit (AU). The AU is defined as the "computational power delivered by a computer executing for one hour at the sustained rate of one GFlop/s", as derived by the top500 list (www.top500.org).

Each researcher applies for accessing to a specific HPC system of one of the seven TA sites, asking for a given amount of CPU core hours on that system, sufficient to realise his scientific project. Once assigned, the core hours are translated in terms of AU for uniformity of global accounting at HPC-Europa level.

Does HPC-Europa 2 have a resource bank?

The resources are allocated to each granted application, on the basis of the request indicated in the application itself. The resources are available for the period necessary to complete the project (in general from the beginning of the visit until some months after the conclusion of the visit).

HPC-Europa 2 does not have a resource bank. But, if a project requires more resources than the ones initially planned, to fully complete the scientific activity, some additional resources can be granted, at discretion of the single sites.

Does HPC-Europa 2 use incentives to encourage use of particular resources?

HPC-Europa 2 does not offer particular incentives for applicants to use specific HPC systems. The selection of the HPC Centres is not only based on the systems available at the centres, but mainly the Applicant asks for a scientific host (from an academic institution near to one of the HPC centres providing TA) to cooperate with, in order to have a scientific cooperation and reach better scientific results.

However it is worth noting that, under the "Common Pool" mechanism, some applications can easily be transferred between centres to maximise the number of users who can benefit from the facilities.

How are resources allocated?

The proposals undergo both a technical evaluation and a scientific peer review.

The technical evaluation is done by a member of staff from the specified TA centre, who will comment on aspects such as the suitability of the facilities requested, the availability of resources requested (such as packages, disk space, compilers and other tools etc), the amount of training likely to be required to enable the applicant to make effective use of the facilities and the feasibility of the draft work plan. This information will be used to provide supplementary information for the selection panel.

The Scientific Users' Selection Panel (SUSP) – composed of 21 eminent scientists representing a broad range of scientific disciplines and independent of any of the facilities - is in charge for the scientific peer review. Each application is revised by two members of the SUSP.

Based on these technical and scientific evaluations, the SUSP meeting makes the final selection of the applications, assigning them to the TA Centres indicated by the Applicants.

Sometimes a given TA centre is not able to offer to the Applicant all the requested resources. If the case the resources must be scaled back, this action is done with reference to both the scientific and the technical evaluations.

The whole application process, from submitting an application through to the delivery of the final report, is based on a web based management tool. Each researcher registers himself and completes the application form electronically. Both the technical evaluation and the host

evaluation forms are completed on line. The SUSP members also complete their pre-selection reviews online in advance of the selection meeting.

A central portal is provided for visitor application forms, supporting documentation (Statements of Support, Technical Evaluation Forms and Host Support Forms), SUSP members' pre-selection reviews and the decisions of the selection panel.

How is feedback collected (about allocations) from users?

Each visitor, during the visit is in contact with the technical staff and the people involved in HPC-Europa at each TA Centre. So each participant is in direct contact with people which can help on the different aspect of the project, including the allocation of the resources.

Furthermore, each granted visitor has the access to a specific protected web site to complete on-line the following documents:

- EU Questionnaire (to provide feedback to the EU on the visit)
- Abstract of the proposal
- Visitor Questionnaire (to provide feedback on the visit, resources allocated, service, training, etc.)
- Host Questionnaire. (completed by the scientific host)
- Short Visit Report (describing the scientific results, to be uploaded once completed the project).

The Visitor Questionnaire represents a formal mechanism for the users to give feedback on their visit and the programme as a whole, and to raise any issues with the relevant TA centre. Furthermore, it provides a QA mechanism for the Transnational Access activity.

Moreover, a User Group Meeting is organised once per year, giving the possibility to the visitors to present their scientific results, but also a further occasion to get in touch again with the representatives of the TA centres.

Does HPC-Europa 2 offer preliminary allocations?

No, HPC-Europa does not offer preliminary allocation. But given the quite high heterogeneity of HPC systems offered in the HPC-Europa Research Infrastructure, the applicants in general already have a preliminary idea of the performances of these systems.

Prospects for further collaboration with PRACE?

Access to the HPC-Europa 2 TA Centres is very important for young researches at the beginning of their computational research activity. During the TA centres visits, they often enhance the scalability of their computational applications and, by means of the cooperation with scientists working in the same research field, they are able to improve the model and the computational methodology adopted. In this way HPC-Europa represent a good step toward the access to more advanced HPC resources as the ones offered by Tier 1 and Tier 0 in PRACE.

6. Nordic Data Grid Facility and the Nordic Resource Trading project

Initial notes gathered during NDGF RT project meeting, 16 November 2011, Espoo, and extended based on project internal reports and discussions with Antti Pursula (NDGF RT Project Manager - CSC) and Pekka Lehtovuori (Development Manager, CSC).

What are NDGF and the Nordic Resource Trading Project?

The Nordic Data Grid Facility (NDGF) is a collaboration between the Nordic countries (Denmark, Finland, Norway, Sweden) established in 2002 to allow researchers in these countries to create and participate to large computational challenges, in particular those resulting from CERN LHC experiments. NDGF is a *production grid* facility that leverages existing, national computational resources and grid infrastructures. It provides and operates the Nordic Tier-1 for WLCG – it is the only distributed Tier-1 of the LHC computing grid.

The Nordic Resource Trading (RT) project was a 1 year project (1.1.2010-31.12.2010) initiated by the NDGF Steering board with the goal of building a prototype system for exchanging IT resources and services within the Nordic countries. It involved four participants: Aalborg University, AAU (Denmark), CSC – IT Center for Science (Finland), University of Oslo, UiO (Norway), and the High Performance Computing Center North, HPC2N, Umeå (Sweden). The project focused on policy and technical issues and delivered an analysis of requirements and scenarios for a future Nordic policy for resource trading, as well as a web portal prototype allowing access to and management of the resources exchanged.

What resources are being/planned to be exchanged?

NDGF does not have a resource exchange system *per se* but does, however, allow access for Virtual Organizations (VO) to a set of resources (mostly computing and storage resources) distributed among its centres.

The RT project investigated various scenarios for establishing a resource trading model that could possibly include a large set of resources such as computing resources, storage capacity, and access to software and expert services.

Where do the resources being/planned to be exchanged come from?

NDGF does not have its own resources. VOs provide resources for sharing from national resource centres and NDGF operates a grid interface for the sharing of these resources together with resource providers.

The RT project examined the possibility of building a Nordic market place where the resources could come from the participating computing centres or the individual users. The most feasible scenario was named “pooled resources”. In this model, each participating country allocates a set of resources to a common pool from where the users (research groups) can select the services they wish to use. Resource exchange does not happen by direct trading between users, centres, or countries. Instead, the resource exchange system sets the framework for international resource access so that overall balance is considered.

What is the unit of exchange?

In NDGF access requests are examined on a case to case basis and no specific unit has been defined to compare the value of the different resources.

Resource trading needs a reliable credit system to facilitate the exchange of resources. The RT project identified the “unit of exchange” question as one of the most difficult issue to tackle when planning to build a resource exchange market, in particular when resources exchanged are of different types (e.g. comparing the value of CPU hours to expert work is not as straightforward as comparing the value of CPU hours on one machine to CPU hours on

another machine). According to the conclusions of the project, the most feasible option seems to measure the value of each service in real currency (e.g. €) and define a conversion factor to produce credit units.

Does NDGF have a resource bank?

NDGF does not have a resource bank.

The "pooled system" envisaged by the RT project does not mention resource banks but the proposed model could make use of such a system.

Does NDGF use incentives?

NDGF does not offer particular incentives for users.

How are resources allocated?

In NDGF, resources are allocated at the VO level.

In the proposed RT pooled model, resources allocation would be a continuous process done at national level. There would be no calls or centralized resource allocation body involved. Each country can grant as much resources exchange credits to national users than what it has committed to the common pool. The national users are supposed to use the credits for accessing services from other countries.

How is feedback collected (about allocations) from users?

NDGF undertakes regular user surveys to get feedback from users.

Does NDGF/RT have preliminary allocations?

N/C

Prospects for further collaboration?

N/C

More information:

http://www.ndgf.org/ndgfweb/home.html

https://resources.ndgf.org/web/guest

7. HP-SEE

Notes from meeting with Ognjen Prnjat, of HP-SEE on 03/02/2011

Ognjen is the project manager of the HP-SEE project and works at GRNET S.A., Greece.

What is HP-SEE?

The HP-SEE project (http://www.hp-see.eu/ "High-Performance Computing Infrastructure for South East Europe's Research Communities") brings together the National HPC infrastructures in the region of South Eastern Europe and the regional Virtual Research Communities of Computational Physics, Computational Chemistry and Life Sciences, aiming at enabling those user communities to get access to HPC resources for their scientific work. HP-SEE, co-funded by DG-INSFO e-Infrastructures unit, is envisaged as a step towards the inclusion of the countries of SEE into the pan-European HPC infrastructures such as PRACE and DEISA.

The countries that participate in the project are: Greece, Bulgaria, Romania, Turkey, Hungary, Serbia, Albania, Bosnia-Herzegovina, FYR of Macedonia, Montenegro, Moldova (Republic of), Armenia, Georgia, Azerbaijan.

The project started in September 2010 and has duration of 2 years.

What resources are being exchanged?

Currently HP-SEE resources include 50 Teraflops distributed in different HPC installations in the participating countries. The biggest being in Bulgaria, where a BlueGene/P system with around 24 Tflop achieved performance is available. The rest of the resources are distributed in Bulgaria, Serbia, Romania and Hungary.

It is envisaged that by the end of the project the total amount of resources available in the infrastructure will be 250 Tflop, enhanced by new purchases in Bulgaria, Romania, Serbia, Hungary and Greece.

Where do the resources being exchanged come from?

The Bulgarian Supercomputer (BlueGene/P) is the national supercomputer in Bulgaria and Belongs to the Bulgarian Supercomputing centre. Resources of this supercomputer are shared between national user communities and the user communities participating in the project.

The rest of the current HPC clusters belong to the project partners or to institutes that are closely related to the partners or with the applications that are supported in the project.

The new resources that will be available in the infrastructure are mainly from new National Supercomputing facilities of the countries that participate in the project.

What is the unit of exchange?

At the moment there is not specialised unit of exchange in the project, and the CPU core hours provided by the accounting systems is used as a rough estimate of resource provision. Resources are shared between partners on a best effort basis and based on the contractual responsibilities of the partners being members of the HP-SEE project.

Applications to be ported and optimized in the infrastructure are mostly pre-selected and stated in the DoW of the HP-SEE project.

Does HP-SEE have a resource bank?

Not at the moment. This is considered as an option for the regional exchange model to be defined within the project.

Does HP-SEE use incentives to encourage use of particular resources?

The incentives are purely technical based on the requirements of the specific applications, the availability of hardware and software suitable for the applications and the expertise of sites that support the applications.

How are resources allocated?

The project has setup an Applications Review Committee (ARC) that consists of experts from the three strategic for the region scientific fields (Computational Chemistry, Computational Physics, and Life Sciences), infrastructure experts and the technical manager of the project. It is the responsibility of the ARC to evaluate the amount of resources that applications require the prioritization of applications, the allocation of resources to the available systems and the approval of new applications that are not specified in the DoW of the HP-SEE project.

The project aims at defining a clear exchange model between the partner countries that will be valid even after the duration of it. Input from similar initiatives is welcome.

How is feedback collected (about allocations) from users?

Feedback is undertaken via mechanisms described in the project. Project reports, surveys, documentation of experience in the wiki, list of publications, helpdesk system are some of the defined way for collecting feedback from the applications.

Does HP-SEE offer preliminary allocations?

Applications are given resources to the so called “home” cluster to test their suitability for porting and optimization to the available in the region HPC resources. The “home” cluster offers limited resources for the above purposes. After that each application sends their request for normal access to the ARC (described above).

Prospects for further collaboration with PRACE?

HP-SEE would like to collaborate with PRACE in the definition of resource interchange models.

HP-SEE is also used as a platform for applications for the region to test their scalability and acquire the technical qualifications for access to the PRACE Tier-0 systems. Further to that HP-SEE would be interested in exploring the possibilities of collaborating with PRACE-1IP WP4 in defining and testing a resource exchange model, as the ultimate goal of HP-SEE is enabling the scientists of the region to have access to pan European research infrastructures for the benefit of their research purposes.

8. LinkSCEEM-2

Notes from email correspondence with Fotis Georgatos and colleagues from LinkSCEEM-2 (February 2011)

Norbert Attig is Head of Applications at JSC (Juelich) and is their representative in the project.

Constantia Alexandrou is the Project Co-ordinator and a Professor of Physics at the University of Cyprus.

Fotis Georgatos is responsible within the LinkSCEEM-2 project for collecting information about user requirements and is presently conducting user surveys.

What is LinkSCEEM?

The major goal of the LinkSCEEM project (FP7 Support Action) is the provision of Tier 1 supercomputing resources to regional user groups in the Eastern Mediterranean (www.linksceem.eu). Only two countries in the Eastern Mediterranean area currently have computers in the Top 500 list, namely Israel and Saudi Arabia.

The objective is to facilitate the creation of an HPC eco-system for the scientists in the region and enable them to integrate their small cluster resources with Teraflops computers coupled with the appropriate training and technical support.

LinkSCEEM has initiated the development of a network of users in various fields who will benefit from the development and integration of HPC facilities. It has carried out surveys aimed at the assessment of regional needs and demands, and organized a number of user meetings for direct contacts with the research communities. A clear need for enhanced e-resources, with associated user support, was demonstrated, while training and educational programs were earmarked as important activities for many scientists in order to be able to use Tera- and Peta- scale computing.

Through the activities undertaken in LinkSCEEM a significant network of scientists has been established in the area. The present project will further develop and expand the scientific links among scientists in the region as well as the collaboration of the Eastern Mediterranean research teams with the corresponding ones in Europe and the USA.

What resources are being exchanged?

All HPC facilities and research institutions of the consortium are currently engaged in user support, training and research activities and therefore this project will provide additional resources in terms of expertise and human resources, in order to enlarge and integrate the user communities in the region. It will promote further collaboration among the partners, integrate computational resources and optimize their usage. In the case of visualization, it will made accessible for the first time advanced and collaborative infrastructure for scientists in the region to pursue activities in digital cultural heritage, presentation of climate data, and synchrotron based research.

Where do the resources being exchanged come from?

Three regional HPC facilities will contribute resources and services in this project. This proposal aims at providing funding for activities that go beyond the services that these facilities offer to their institutional and national users, and are aimed at engaging virtual research communities at the regional scale.

LinkSCEEM-2 partners include three High Performance Computing (HPC) centres from the Eastern Mediterranean region who have committed a significant portion of the computational, data storage and visualization resources to the project. The three HPC centres are:

- Computation-based Science and Technology Research Center (CaSToRC, Cyprus)
- National Authority for Remote Sensing and Space Sciences (NARSS, Egypt)
- Bibliotheca Alexandrina (BA, Egypt)

The project foresees the engagement of additional resource provider organizations from the region in order to facilitate the emergence of HPC ecosystem in the region.

There are, however, currently no local national Tier 1 resources available to European researchers in eastern Mediterranean countries, so there are no Tier 1 resources to allocate. At CaSToRC (Nicosia, Cyprus) they are in the process of procuring a 40 TF cluster system. (The tender closes in March 2011) At the Egyptian institutions Bibliotheca Alexandrina (BA) and National Authority for Remote Sensing and Space Sciences (NARSS) there exists a cluster and a one rack BG/L which currently cannot be used from remote institutions. The project partners are working hard on these bottlenecks to overcome them. They will optimally integrate resources by linking established lead HPC centres (namely Juelich Supercomputing Centre (FZJ-JSC) and of the National Center of Supercomputing Applications (NCSA)), developing and sharing best practices for managing these resources; The aim is to ensure both the embedding of the regional HPC facilities into the European HPC eco-system and the development of an international dimension.

What is the unit of exchange?

LinkSCEEM is currently implementing the preparatory access procedure: http://www.linksceem.eu/hpcaccess

It is written around the concept of core-hours without any discrimination between architectures. THIS MAY CHANGE AT ANYTIME. Of course, not all resources are indeed the same. This appears intentional since in the first round we need to collect experience amd want to keep things as simple as possible for the regional users and sites.

LinkSCEEM would be keen to understand the requirements for exchange of resources in PRACE to maximise collaborative and participation opportunities. Indeed it would be good to include some ideas about the requirements for exchange of resources in the PRACE report.

Does LINKSCEEM use incentives?

The core * hours product might be adequate to drive users towards the resources which provide for them the best service/products. This has to be investigated.

How are resources allocated?

A resource allocation mechanism has been developed and will be applied to upcoming users; a resource exchange procedure has not been developed yet. Maybe that at a later point in time a resource exchange may be an opportunity for LinkSCEEM users.

Of prime importance is the set up of a project Resource Allocation Committee administering user requests for access to partner resources

How is feedback collected (about allocations) from users?

Not yet applicable, however information is being collected about requirements from user surveys and also via the Winter Schools of LinkSCEEM and PRACE which are being held in Cypus.

Does LinkSCEEM have preliminary allocations?

Yes, we call them Preparatory access; the first draft CfP is being documented:

http://www.linksceem.eu/hpcaccess

Prospects for further collaboration?

LinkSCEEM is keen to understand the requirements for exchange of resources in PRACE to maximise collaborative and participation opportunities. In addition, the project will contribute to the integration of compute resources with the data and visualization resources that are required for applications to have optimal scientific impact. Moreover this project will contribute to the sharing and coordination of methodologies and practices, catalysed by the presence of Tera-scale systems and by the transfer of high-level expertise and know-how from internationally reputed institutions. This in turn will create greater integration between the resources that are actually involved in the project and the local clusters operated by research groups in the region, which will favour the development in the region of a HPC eco-system, similar in spirit to that which PRACE is promoting in Europe, but at a more modest scale.