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

D4.3.2
Cross-National Programme for Tier-1 Access Pilots

Final

Version: 1.0
Author(s): Petri Nikunen, CSC and Alison Kennedy, EPCC
Date: 26.06.2012

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: D4.3.2	
	Deliverable Nature: Report	
	Deliverable Level: PU *	**Contractual Date of Delivery:** 30 / June / 2012
		Actual Date of Delivery: 30 / June / 2012
	EC Project Officer: Thomas Reibe	

* - 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 Tier-1 Access Pilots	
	ID: D4.3.2	
	Version: 1.0	**Status:** Final
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2003	
	File(s): D4.3.2.docx	
Authorship	**Written by:**	Petri Nikunen, CSC and Alison Kennedy, EPCC
	Contributors:	Pedro Alberto, UC-LCA; Andrew Emerson, CINECA; Juha Fagerholm, CSC; Chris Johnson, EPCC; George Karagiannopoulos, GRNET; Miroslaw Kupczyk, PSNC; Ioannis Liabotis, GRNET; Georgi Prangov, NCSA; Wim Rijks, SARA; Stéphane Requena, GENCI
	Reviewed by:	Sergi Girona, BSC; Thomas Eickermann, FZJ
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	15.05.2012	Draft	Initial version
0.2	23.05.2012	Draft	Revisions by Wim
0.3	24.05.2012	Draft	Revisions by Juha
0.4	29.05.2012	Draft	Revisions by Ioannis and George

0.5	31.05.2012	Draft	Revisions by Andrew, Pedro, Chris, Georgi, Wim, and Stéphane
0.6	07.06.2012	Draft	Revisions by Alison
0.7	08.06.2012	Draft	Revisions and additions by Alison (including comments from Wim)
0.8	11.06.2012	Draft	For internal review
1.0	26.06.2012	Final	

Document Keywords

Keywords:	PRACE, HPC, Research Infrastructure, DECI, Resource exchange

Table of Contents

List of Figures

List of Tables

References and Applicable Documents

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

[2] D4.3.1: *Cross National Programme for Access to Tier-1 Resources*, PRACE 1IP, http://www.prace-ri.eu/IMG/pdf/D4-3-1_1ip.pdf (2011)

[3] D2.1: *Migration from DEISA2 to PRACE-2IP*, PRACE-2IP, http://www.prace-ri.eu/IMG/pdf/d2.1_2ip.pdf (2011)

[4] *A Strategic Agenda for European Leadership in Supercomputing: HPC 2020*, IDC Final Report of the HPC Study for the DG Information Society of the European Commission http://www.hpcuserforum.com/EU/downloads/SR03S10.15.2010.pdf (2010)

[5] *Financing a Software Infrastructure for Highly Parallelised Codes*, IDC Final Report for the DG Information Society of the European Commission, http://cordis.europa.eu/fp7/ict/e-infrastructure/docs/fisofi4hpc.pdf (2011)

[6] D5.1.2: *Requirements for Industrial Users*, PRACE-1IP, Not public (2011)

List of Acronyms and Abbreviations

AISBL	Association Internationale Sans But Lucratif (International Non-Profit Organization)
BSCW	Be Smart – Cooperate Worldwide; a web-based system that offers shared workspaces
CINECA	Consorzio Interuniversitario, the largest Italian computing centre
CPU	Central Processing Unit
CSC	IT Centre for Science (Finland)
DAAC	DECI Application Acceptance Committee
DART	Distributed Accounting Report Tool; a Java web-start application to retrieve machine accounting data
DECI	Distributed European Computing Initiative; a scheme through which European computational scientists can apply for single-project access to Tier-1 resources through PRACE
DEISA	Distributed European Infrastructure for Supercomputing Applications; EU project by leading national HPC centres. Ended in 2011
DG	Directorate General (department of the European Commission)
DPMDB	DECI Project Management Database; a web-based application to view and edit details of DECI applications and projects
EPCC	Edinburg Parallel Computing Centre (represented in PRACE by EPSRC, United Kingdom)
GHz	Giga (= 10^9) Hertz, frequency =10^9 periods or clock cycles per second
GPGPU	General Purpose GPU
GPU	Graphic Processing Unit
GRNET	Greek Research and Technology Network
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing
HPC-Europa	An EC-funded programme by which scientists can carry out short research visits to collaborate with a research department working in a similar field
IBM	Formerly known as International Business Machines; a multi-disciplinary company that develops and manufactures supercomputer technologies
IDC	International Data Corporation; a market research and analysis firm specialising in information technology, telecommunications and consumer technology
ISV	Independent Software Vendor
LRZ	Leibniz-Rechenzentrum (Germany)
MS	Microsoft; an international software company that develops and licenses software products for consumers, businesses and public sector organisations
NCSA	National Centre for Supercomputing Applications (Bulgaria)
PDF	Portable Document Format
PI	Principal Investigator
PRACE	Partnership for Advanced Computing in Europe; project acronym
PRACE-1IP	First Implementation Phase of PRACE
PRACE-2IP	Second Implementation Phase of PRACE
PRACE-3IP	Third Implementation Phase of PRACE
PRACE-RI	PRACE Research Infrastructure
PSNC	Poznan Supercomputing and Networking Centre (Poland)

R&D	Research and Development
RZG	Rechenzentrum Garching (Germany)
SARA	Stichting Academisch Rekencentrum Amsterdam (the Netherlands)
SE	Scientific Evaluation
SSC	Scientific Steering Committee (of PRACE)
SUSP	Scientific Users Selection Panel of HPC-Europa
TE	Technical Evaluation
TeraGrid	An infrastructure for open scientific research in the USA; ended in 2011
Tier-0	HPC systems hosted by the PRACE RI; the largest European systems
Tier-1	National or topical HPC systems
Tier-2	Regional or campus HPC systems
UC-LCA	Laboratory for Advanced Computing in the University of Coimbra (Portugal)
WP	Work Package
XE	Cray's supercomputer technology
XSEDE	The Extreme Science and Engineering Discovery Environment; Follow-up of TeraGrid
XT	Cray's supercomputer technology

Executive Summary

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.

The project envisaged interchange of cycles being used for three core activities. Firstly, to provide a pool of resources for peer-reviewed single-project access via DECI (launched under DEISA and continued in PRACE), secondly to provide longer term access to Tier-1 for science communities (complementary to Tier-0 programme access) and thirdly to provide resources for non-proprietary research and code development, testing, scaling and optimisation for European industry (again to complement Tier-0 provision). In each case, the activity is intended to offer an alternative to Tier-0 and also to act as a ramp to Tier-0, for those users requiring access to European facilities (but not at the Petascale level).

As its initial starting position, PRACE planned to take account, wherever appropriate, of the experiences and good practices developed by the DEISA and DEISA2 projects, whilst also taking account of the changing role of national systems providers as Tier-1 in the European HPC ecosystem. The rapid expansion of HPC provision in Europe, at both of Tier-0 and of Tier-2, gives Tier-1 a clear role as the intermediary to facilitate movement between "campus" or "regional" facilities offered by Tier-2 and European "leadership class" computers offered by Tier-0.

This document evaluates the outcome of the DECI pilot undertaken in PRACE (including a review of processes and practices). It also summarises the investigations into the possible future role of Tier-1 in science community access and in industry access.

While the DECI processes and practices are currently judged to be "fit for purpose", this document suggests that they could be improved, in particular to provide a more uniform look and feel between Tier-0 and Tier-1 and to automate more of the processes to free up staff time for user-facing activities. This is particularly pressing in the light of the continuing expansion of the number of Tier-1 facilities which introduces increased complexity to the management of DECI.

This document also recommends that consideration is given to running pilots for both science community access and industrial access to Tier-1 and to the delegation of responsibility for Tier-1 to a PRACE Optional Programme, with a remit to plan and develop Tier-1 strategy over a longer timeframe and to manage and administer the increasingly wide portfolio of Tier-1 activities.

These ideas will be explored further in WP2 of the PRACE-2IP project.

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. As its initial starting position, PRACE planned to take account of the experiences and good practices developed by the DEISA and DEISA2 projects for the DECI (DEISA Extreme Computing Initiative), which is continuing in PRACE as the Distributed European Computing Initiative. DECI was launched under the EU funded project DEISA in 2005 to support pan-European computational science research. Through DECI, any European research group could apply for computing resources, and if its research proposal was considered good enough, it was accepted, subject to sufficient resources being available for allocation. Resources were given to projects at no cost to them.

The experiences of, amongst others, HPC-Europa and HPC-Europa2, the European Commission's Transnational Access and Mobility programme in HPC, were also taken into consideration by PRACE.

When the task was originally proposed, the intention was to collect and analyse information on which to make informed decisions about a future, but as yet undefined, Tier-1 resource exchange programme. However, one of the European Commission's objectives for the PRACE-2IP project was then 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. Hence an early decision to continue DECI with the same acronym (but changed meaning) in PRACE was taken. DECI was migrated to PRACE in 2011. The processes of DECI were migrated nearly "as such". The processes evolved to their current state during DEISA and DEISA2. The processes were documented in the deliverable D2.1 of PRACE-2IP in 2011 [3]. Until this date, the processes had not been evaluated or alternatives had not been actively sought.

With the advent of Tier-0, the European-level "leadership class" computers, the role of Tier-1 has begun to change. The first deliverable of this task, entitled "Cross National Programme for Access to Tier-1 Resources" [2] presented resource interchange concepts to PRACE members and outlined how a DECI pilot call would be run.

This deliverable analyses the results from the pilot call and looks at how the concepts can be further developed to provide a model for European resource exchange at the Tier-1 level for both industry and for science communities, in addition to the existing provision for single-project access. It also looks at the changing role of single project Tier-1 access and the possible role of an optional programme in managing Tier-1 resources and the further evolution of Tier-1 provision.

The purpose of this document is therefore to evaluate existing Tier-1 processes, to assess their suitability in the short and medium term (given the continual evolution of PRACE) and to suggest possible alternatives to them in the light of further evolution. The intention is to make assessments of possible future models, taking into account ease of use for both end-users and PRACE partners, including stronger harmonisation of or co-ordination between Tier-1 and Tier-0. However, any proposed or recommended solutions must, of course, be acceptable both to the national agencies or centres who offer the Tier-1 resources to PRACE and to the governing bodies of the PRACE RI, under whose aegis the Tier-1 access programmes are run.

The document is primarily intended for the Council of PRACE AISBL, but it also is of interest for all PRACE-2IP personnel involved in Tier-1 and to the Tier-1 systems governance bodies.

2 Role of Tier-1 Access in PRACE

The following sub-sections are short introductions explaining the role of the different types of Tier-1 access (either existing or proposed) in the PRACE infrastructure.

2.1 Single Project Access (DECI)

DECI (Distributed European Computing Initiative) provides a resource exchange programme for Tier-1 resources within Europe. The programme has been running for many years, firstly under DEISA and since September 2011 under PRACE. Participating DECI countries provide resources of around 5% of their national resource to DECI. These resources are then distributed across successful DECI projects with the aim that projects within each country receive an amount of computing resources roughly equal (quantitatively) to the amount contributed but qualitatively different (i.e. on different classes of architecture). Up to 15% of the resources contributed by participating countries are allocated to "external" projects, i.e. those from countries not contributing resources themselves.

Projects are submitted by PIs and reviewed first technically by PRACE staff and then scientifically by national peer review panels. For projects received from countries not offering Tier-1 resources, the scientific reviews are performed by the HPC-Europa2 "SUSP", a panel of 21 pan-European scientific experts whose expertise is drawn from a broad spectrum of scientific disciplines. The peer-review is therefore undertaken in a distributed manner by delegating authority to existing expert European peer review committees, working to common PRACE peer review guidelines.

Successful projects are given access to Tier-1 resources for a period of 12 months from an initially defined starting point. As in Tier-0 level, in DECI calls users will benefit from a greater diversity of available resources and machine configuration throughout Europe including Cray (XT and XE), IBM Blue Gene/P and Blue Gene/Q, IBM Power6, a range of large clusters (various processor and memory configurations and some other new architectures) and hybrid systems (clusters with GPGPU accelerators).

The original motivation for CPU resource exchange at the European level, when introduced by the DEISA project in 2006, was to extend the variety of architectures and machine configurations available to European computational scientists by offering peer-reviewed, single-project access to the large national systems of the partners in the DEISA consortium, based on scientific need, rather than country of work. For example, a QCD researcher based in a country where the national system was an IBM Power4 could apply for time on a Blue Gene/L which better suited his or her needs. Conversely, a scientist based in a country where the national system was a Blue Gene/L, whose code was better suited to a system with a larger memory per node could apply via DECI for access to an alternative architecture. Each partner contributed a minimum of 5% of the cycles of their national HPC resources to a pool of CPU. The national systems on which CPU was allocated were "Top 100" machines (among the largest HPC facilities in the world). Proposals were peer-reviewed by the national committees.

Initially, the applicants to DECI were typically individual researchers or members of small consortia from the same country and the main motivation for applying was to obtain access to alternative architectures, plus applications enabling support from DEISA so that the code could be optimised for the architecture to which it was ported. However, over the years, the average consortium size became larger and typically included members from several countries, in line with the increasing globalisation of scientific research. The emphasis for DECI therefore gradually shifted from resource exchange to resource pooling. It became

increasingly difficult to label a project or consortium as, for example, "German" or "British" (purely on the basis of the nationality of the PI) when the multiple investigators came from different countries. The introduction of PRACE Project Access also changed the landscape, in that European scientists had a route by which they could apply for larger allocations on the very biggest European machines. These machines were no longer available via DECI. The amount of resource now available via DECI, although continuing to grow, is composed of fractions of a much larger number of smaller machines (around 20 as opposed to around 8 in the early days of DECI).

The current situation is that there are a large number of Tier-1 partners, whose national systems range from being quite modest in size to being close in size to the Tier-0 systems. Most proposals are submitted by consortia which have members from several different countries. The challenge for the future is to develop a resource pooling model which takes account of the changing position of DECI (as a stepping stone to Tier-0 for many consortia) and of the increasing internationalisation of research.

2.2 Community Access

2.2.1 *Principles for Scientific Community Access*

The reasoning behind the recommendation to offer scientific communities access is similar to the envisaged programme access in PRACE for Tier-0 resources, to let the scientific communities decide themselves the final allocation of resources to sub-projects or individuals and to have access to resources longer than in regular calls. Structured scientific communities in some areas have already been established and they could serve as a basis for this type of access.

2.2.2 *DEISA2 Virtual Communities*

The community access was already started in the framework of DEISA where they were called virtual communities. In DEISA2, there were five selected virtual communities:

- EFDA (European Fusion Development Agreement)
- EUFORIA (EU Fusion fOR Iter Applications)
- ENES (European Network for Earth Systems Modelling)
- LFI-PLANCK (PLANCK Low-Frequency Instrument community)
- VPH (Virtual Physiological Human initiative)

The virtual communities had a long term access to DEISA resources, and the yearly share of resources was decided on the project management level. They were allocated 15% of all CPU resources. The exact amount assigned for each community was based on virtual communities' needs and their application; the final amount was then negotiated by DEISA with the project leader.

Virtual community support is part of the PRACE-2IP work plan. Some of the virtual communities continued in the period between DEISA2 and PRACE-2IP using additional resources provided by the former DEISA2 partners. The virtual communities that continued also used the basic infrastructure (ex-DEISA network etc.) provided by PRACE-2IP.

2.2.3 *PRACE Programme Access*

The PRACE Scientific Steering Committee (SSC) launched a call for Expressions of Interest (EoIs) for Programme Access in October 2011. The results were discussed in a meeting at the end of November 2011. The analysis of the EoI was "the SSC recommend not to give the resources to the communities for own distribution, as the access committee is responsible for". From the EoI, the SSC understand that a multi-year access for specific projects may be adequate, and a pilot/test call was established. Based on the results, a call for multi-year access was launched by PRACE AISBL for Tier-0 resources in May 2012. This is very different from Programme Access: resources are not managed by a community.

A few of the DEISA virtual communities expressed their intent to continue after DEISA2, and they also participated in the EoI call for Programme Access in 2011. As the SSC stated in the summary of the call for EoIs, several communities asked for an access to both Tier 0 and Tier 1 resources.

Some of these Scientific Communities (e.g. PLANCK) have participated in the PRACE DECI calls for continued support of European Tier-1 resources and dedicated network access.

Based on experiences from DEISA, PRACE could provide also Tier-1 resources for scientific communities. No decision for it has been made yet, because it was agreed to wait for the result of the call for EoI for Tier-0 Programme Access.

2.3 Industrial Access

Since its inception PRACE has regarded the mission of setting up a world class European HPC service not only as fundamental tool for top-notch scientific research but also as an enabler for European industry competitiveness. In the latest strategic PRACE AISBL discussions, industry is seen as the second pillar of PRACE mission, embodying not just industry access to CPU cycles but also to HPC competence.

Up to now industry access to PRACE resources has been limited to some pilot projects running on Tier-0 systems for open R&D. The work package WP5 of PRACE-1IP has been concerned with relationships with industry, namely in which ways can industry access Tier-0 resources. Issues like confidentiality, ISV access and commercial vs. open R&D usage have been discussed and debated. PRACE-2IP WP9 is currently setting up a call for enabling applications for industry.

The industry seminars organized by PRACE have shown that in Europe there is a genuine industry interest for HPC. In particular, the 2011 seminar in Stockholm, which had a special focus on small and medium-sized enterprises, showed that while they have in general small requirements in terms of core count for running their applications, they often lack competence in either porting or using parallel applications. Most of the times it does make sense to these companies to invest in a parallel computer, however small, and also hire personnel to manage it. But, even if they do, in most cases they lack the HPC competences of major supercomputer centres to run the codes effectively. All of this builds the case for developing a model for PRACE Tier-1 industry access, including not only CPU cycles but also the competence for adapting, scaling and running industry application codes. Naturally, there should be a continuing dialogue with industry representatives in order to assess their needs and evolving requirements. The Industrial Advisory Committee being set up within PRACE AISBL and the PRACE industry seminars seem the natural places to do that.

2.4 Preparatory Access

Similarly to PRACE Preparatory Access for Tier-0 resources, which is done through rolling calls, opened all year round with cut-off days, this type of access could be opened also to Tier-1 resources. It could be opened both for academic and industrial users. Every single project would be submitted to a technical evaluation every 3 months at defined cut-off dates. Additionally, limits for CPU time may be applied as in Tier-0 and Tier-1 regular calls (for instance a maximum 2% of the total amount of contributed CPU time) to preparatory access.

The *Preparatory access* calls on Tier-1 systems would give researchers the possibility to apply for:

- access to perform code scalability tests and code optimisations, and
- support for code development and optimisation with the help of software experts amongst the PRACE project partners.

The purpose of *Preparatory access* would be to allow researchers to optimise and test application codes before applying to the regular *Project access* type of Tier-0 or Tier-1 calls.

In the PRACE Preparatory Access for Tier-0 resources, there are three types of access. The same types could be adopted for Tier-1 preparatory access.[1]

Type A: Code scalability tests

Applicants may use this *preparatory access* to obtain scalability plots of the performance of the code on a certain HPC system. This allows the applicant to obtain up-to-date information when applying to PRACE *Project access* type of calls. The maximum allocation period for type A is 2 months which includes the submission of the final report.

Type B. Code development and optimisation

Applicants may use *preparatory access* to undertake adaptations to the application code. Applicants will need to describe the proposed development plan in detail, including the human resources (and expertise) available to implement the project. The maximum allocation period for type B is 6 months which includes the submission of the final report.

Type C. Code development and optimisation with the support of PRACE experts

Applicants use this type of access for the same purposes as type B. In addition, under *preparatory access* type C, the applicants may require the support of PRACE experts to improve the application code. The maximum allocation period is 6 months which includes the submission of the final report. This has been piloted as part of PRACE-2IP and will be further elaborated in PRACE-3IP.

The award of *preparatory access* proposals of types A, B and C is limited by the amount of resources available. For type C, there are also limitations regarding the availability of the support from experts (in terms of staff hours) and the type of expertise required (the expertise available is not the same in all fields of computer science, applied mathematics and those of the application science or engineering field).

Preparatory access would allow applicants to access PRACE Tier-1 systems for testing and/or adapting computer codes in preparation for the participation in *DECI access* calls. Production runs would not be allowed as part of preparatory access. Preparatory access applications would be evaluated based on their merits on technical grounds. There would be no evaluation on scientific merit.

[1] Tier-1 preparatory access does not currently exist. Applications enabling assistance is provided as an integral part of DECI, along with CPU, however this model could be reviewed in future.

3 Valuation of Resources

3.1 Policy in DEISA DECI and PRACE Pilot calls

In DEISA, the only resource available was the compute resource. To compare compute resources on different platforms the "DEISA standard core-hour" was introduced. This was defined as a wall clock hour on an IBM Power4+ (1.7GHz) core. Resources from all contributing machines were converted to the "DEISA standard core hour" according to pre-established conversion factors, based on running a number of applications from the DEISA Benchmark Suite and comparing the performance to that on the IBM Power4+. For different processors belonging to the same family, for instance Intel Xeon or Nehalem running on different clock speeds, sometimes an estimate was made based on this difference in clock speed.

Resource requests in DECI proposals were all converted to this common standard core-hour.

This approach has a number of drawbacks:

1) Converting a DECI resource request into standard core-hours required some effort, because it had to be established how the Principal Investigator (PI) came to his estimate of the required resources. It often appeared to be based on test runs on a platform not present in the DEISA infrastructure. This led to the additional problem of determining a conversion factor for these platforms.
2) With increasing number of sites that participate in the DECI Tier-1 calls, the number of architectures that will contribute to the infrastructure will increase and this means that the process of determining a conversion factor will have to be performed more often.
3) The use of the Power4+ processor as standard core only works as long as the benchmark suite remains static. If we have to replace codes or add new ones we get into trouble, because the standard Power4+ core is no longer available within PRACE, so it is not possible to run these applications on a Power4+ to obtain the conversion factor directly. Instead one can run the new application codes on a processor that has a well established conversion factor and use these results as a reference to determine a conversion factor for the new system. This approach could also be used to choose a new standard-core processor.
4) It might even be argued that a common conversion factor hasn't any predictive value on the efficiency of a particular code on a particular architecture.

Moreover, application benchmark based conversion factors have two generic problems:

1) They depend on the choice of applications. Platforms cannot be benchmarked against all applications of a benchmark suite, because some platforms would get too poor conversion factors that way. Some might say that it reveals the true performance of a platform, but it doesn't tell the whole story. The platforms should be benchmarked against applications they are suitable for. Because there are no other ways to determine the suitability of applications for platforms, except benchmarks, conversion factors based on them are unreliable.
2) As the conversion factors are based on ratios of benchmark results on two platforms, they underestimate the value of fast processors. For an investigator it is in general

much better to have 1 million core hours on a fast processor than 10 million core hours on a processor whose speed is one tenth of the fast one.

3.2 Future Challenges

3.2.1 *New Technologies e.g. GPUs*

The recent availability of clusters with GPUs for Tier-1 presents a serious challenge for defining a resource exchange model. Unlike the situation for computer hardware based on general commodity chipsets (Intel-like, PowerPC etc.), where performances between different processors usually vary only by a small factor, for applications which have been written or modified to exploit GPUs the increase in performance can be very many times that of the equivalent non-GPU or unmodified version. In fact, it is not uncommon to have speedups of the order of ten, twenty or even in some cases, over a hundred times compared to the versions running on non-GPU hardware. These speedups should be compared to the resource exchange models employed within the DECI calls of DEISA or PRACE Tier-1 where conversion factors between different architectures are usually smaller than two2. The lack of this consideration in the resource exchange model penalises both: 1) users who have appropriately enabled code but no access to GPU resources and 2) sites hosting the resources that are not compensated for the extra hardware costs and possibly also the extra application enabling efforts required. However, it is not straightforward to assign a single scaling factor for GPU clusters for a number of reasons:

- GPUs are fundamentally different to conventional processors used for HPC so it is not possible, for example, to compare clock speeds in order to have an approximate value of the speedup possible with a GPU;
- To date only a few applications have been modified to exploit GPUs (up to now almost exclusively with NVIDIA's CUDA library) , and often GPU acceleration has been applied only to certain parts of the code;
- There is a wide variation of reported speedups compared to non-modified codes ranging from factors of 2-3, to factors of 10, 100 or more (as mentioned above);
- Some Tier-1 sites have GPUs only on a subset of the available nodes;
- It is not trivial for accounting procedures to gauge how much time has been spent on a GPU during a batch job. For this reason many sites consider the CPUs and GPUs on a node as a single resource and users get charged in "node-hours" (see below).

A further complicating issue is that within PRACE there are different classes of GPUs, e.g. in the NVIDIA range M2070, M2070Q, M2090, as well as hardware from other vendors. Before bothering about these issues it is worthwhile to consider the policies applied by Tier-0 sites, where a similar discussion has already taken place, and policies that sites apply to their local or other non-PRACE users.

GPU resources definition in Tier-0 and sites local policies for non-PRACE users

[2] One of the largest conversion factors between two different architectures is 9, which is the conversion between Power6 and PowerPC of Blue Gene/P.

Tier-0 users and those users requesting time directly from computer sites can be considered together since compute time is usually awarded directly on a site, rather than through the Tier-1 resource pool.

- The majority of GPUs for *Tier-0* reside on Curie's hybrid partition where each node contains two GPUs and two CPUs. Resources are allocated as "atomic units", i.e. as whole nodes, so it is not possible to request for example only CPUs and users are charged in "GPU-hours". To be allocated on the hybrid partition, applications must be able to exploit the architecture according to a set of minimum requirements.
- For sites offering GPUs to their own national users, policies are generally similar to those adopted by Tier-0, i.e. users are charged in "node-hours" according to the nodes allocated to the job. Some sites offer different types of architectures to their users: in the case of CINECA both Power6 and hybrid CPU/GPU resources are offered but no scaling factor is currently applied to resources on the different systems, i.e. 1 core hour of Power6 time is considered the same as 1 hour of hybrid cluster time. Because of the different clock rates of the main compute processors (4.7Ghz for Power6, 2.4 Ghz Intel Westmere) this amounts to a scaling factor, but was chosen to simplify accounting procedures rather than as a measure of possible performance gain.

Although considering these cases is useful they don't actually help formulating a new resource exchange model because if one applies for only one resource there is no need to apply a conversion factor.

3.2.2 *Storage and Licences*

Research projects need storage for different purposes:

- Small amount of backed-up storage for simulation inputs, source codes, and personal documents.
- Large amount of backed-up storage for archiving simulation results.
- Space for temporary simulation data.

Licences in turn are required by software providers. There are different licensing models:

- License is owned by a user. The most usual case is that the software developer allows a computing centre to install their software on the computing centre's servers, but wants to limit the usage somehow. For example users need to accept some rules or pay money. The computing centre makes sure that software developer's restrictions are enforced.
- License is owned by a computing centre. Typically the computing centre has purchased from the code owner a licence for a certain amount of instantaneous processes of software. When the software is started on the computing centre, it connects via the internet to the computing centre's licence manager, which either allows or disallows the new process. For example, compilers on computing servers work this way. A generalisation to this is software packages that are run with the computing centre's licence on a user's institution's computer or even on her/his own workstation.

Currently storage and licences are provided implicitly to a project as a supplement of CPU time. If one or both are unavailable at a site, it makes this site unsuitable for the project. Resources that do not need to be available at the execution site (e.g. archive space) can be offered by the project's home site if execution sites are not willing or able to provide them.

To reward sites that provide storage and licences, there should be a mechanism to evaluate their worth.

3.2.3 *Exchange with External Organisations*

European research is not confined to European borders. Many European research projects have non-European collaborators, and vice versa. And many European research projects would like to use non-European resources. For these reasons, it might be advantageous to exchange resources with other organisations, for example XSEDE in the USA.

In the resource exchange model, external organisations could be considered analogous to PRACE sites: External organisations provide resources to PRACE and receive resources back for their own research projects. This would give an opportunity to exchange resources between PRACE and external organisations using the juste retour principle (without contributions for virtual communities or external European research projects).

Before exchanging resources with other organisations, PRACE and these organisations should agree on the currency with which the resources are exchanged. It would be best (for PRACE) if external organisations could accept the currency that PRACE uses.

3.3 Development Ideas

The growing importance of hybrid resources such as GPUs demands a re-examination of the resource exchange mechanism. The drawbacks of the current mechanism are obvious even without GPUs and they have become more apparent as GPUs have come into use. Other issues that call for changes are inclusion of storage and licences in the mechanism, and exchange with external organisations.

It is not clear how to proceed on including hybrid resources such as GPUs in a resource exchange mechanism. One idea could be to adopt a benchmark suite of GPU-enabled codes, and then to extract an average scaling factor (this is in fact one of the main uses of benchmark suites). But there may be problems with this approach:

- The number of codes available for such a suite is limited since until now only a few have been enabled for GPUs. This number will increase with time but will still be small since most applications will require modification and some will never be suitable for acceleration;
- Unlike with conventional codes, speedups with GPUs vary much more (as mentioned above). An average speed-up value will therefore depend strongly on which codes are present in the suite;

Two alternatives have been proposed. In the first alternative, a table of conversion factors would be provided for each individual application. For the well-known third party codes, this could be done in advance or even extracted from previous benchmark runs. However for home grown user codes this would imply running the codes on all platforms during the technical evaluation phase. This would put additional pressure on that process, and would relieve the drawbacks of the current resource exchange mechanism only partially. A variation

on this approach could be to charge on the capacity of the CPUs on the node, multiplied by a minimum speed-up factor, specified by the hosting site.

In the second alternative, the conversion factors would be related to monetary value of the resources. It would allow exchange of CPU time, storage, and licences in an easy way. No money would be paid by anybody, but commitments and assignments would be done in Euros. Monetary value based conversion factors might be easiest to accept also by external organisations. The very purpose of conversion factors is to prevent exchanging cheap resources with more expensive ones, so why not using monetary value as a criterion? On Tier-0 level, the costs are assessed by an external auditor that provides cost per CPU hour figures for each machine. The final figures are made public, but the internal cost structures are not exposed to anyone. This kind of approach might work well also for Tier-1.

Our conclusion is that the ideal practice would combine the two proposed alternatives. The first alternative would be used to convert resources stated in the DECI proposals to Power4+ core hours, or to new standard processor core hours. In this process it is important that the method is benchmark based. The second alternative would be used to exchange resources between sites. In this process it is better if the method is monetary value based, although sites providing GPUs only get a modest return for their contribution. The processes to valuate the resources will be analysed further in the optional programme (for details see Sec. A).

4 Review of DECI Pilot Processes

4.1 Call for Proposals

The first DECI call of PRACE (DECI-7) was synchronised with the 3rd PRACE Project Access Call to allow applicants to apply for Tier-0 and Tier-1 projects at the same time. The application forms were separate however, so the applicants needed to choose which call they answered. No re-routings of proposals were made between Tier-0 and Tier-1, so if a Tier-0 proposal requested too few resources or a Tier-1 proposal requested too many resources, they were simply rejected or cut back, regardless of the scientific rating of the proposals.

4.2 Call Launch and Publicity

The synchronised Tier-0 and Tier-1 calls (DECI-7 and DECI-8) were advertised with a single call text which was published on www.prace-ri.eu by the PRACE press team. The call was also disseminated to already existing contacts of the press team, with language translations being done by some of the PRACE partners to publicise the call locally. In DECI-7, the call text was prepared mainly by PRACE AISBL with input from WP4 of PRACE-1IP. In DECI-8, on the other hand, the call text was produced jointly by Tier-0 and Tier-1 representatives, before being approved by PRACE AISBL.

User feedback, as collected by Tier-0 and Tier-1 representatives, suggests that many applicants have difficulty filling in the application forms. These difficulties not only include uncertainties about how many hours to ask for, memory requirements of the applications to be used, etc., but also the inclusion in forms of PRACE-specific concepts such as Tier-0, Tier-1, standard hours, among others. These issues will be addressed more deeply in the call texts and in the documentation supplied for guiding applicants, as this information is key in the technical evaluation and in the assignment of resources to projects.

4.3 Proposal Submission Mechanism

The proposal submission mechanism used in the PRACE DECI pilot calls is virtually unchanged to that used in DEISA. Applications are submitted via Word documents to a specified email address which are then uploaded onto the PRACE BSCW server. Various details of the proposals are also stored in the DPMDB (DECI Project Management DataBase) by PRACE staff using copy-paste approach. This is clearly time-consuming and error-prone.

The recommendation is to change to an online submission mechanism. There are many arguments for this: for users it requires the same effort to fill in a web based or Word based form, but an online submission mechanism makes it much easier to track the subsequent processes. For PRACE staff an online submission process is much easier and more reliable, as it avoids the copy-pasting between Word files and the DPMDB. This impacts on many tasks in the process since the DPMDB is also used to send out confirmation or rejection letters for example. An online mechanism would also avoid unnecessary email, spam etc., and it could integrate a help or manual to give advice on how to fill in the form. In addition, it would allow the proposal submission to be integrated with the technical and scientific evaluations.

Bearing in mind the requirements of the applications system, two online submission tools are being considered for DECI: that used by Tier-0 and the one employed by HPC-Europa2. From a point of view of integration of the Tier-0 and Tier-1 process the suggestion would be to adopt the Tier-0 version, which should not require many modifications for Tier-1. This tool would also emphasise to users the possible transition from Tier-1 to Tier-0. It would be

necessary however to integrate the submission tool with the DPMDB. A taskforce has been formed to evaluate the Tier-0 and the HPC-Europa2 and come up with the requirements for the online submission system.

4.4 Technical Review of Proposals

Statement - Problems

After the submission deadline the DECI proposals are evaluated first technically and then scientifically. Technical evaluation is part of the project review process and is determining the technical feasibility of the proposed projects. Transparency and fairness to the users must be ensured in the process. Technical evaluation is taken into consideration by the peer review panel for the ranking of the applications.

The technical evaluation is currently based on the technical evaluation form that is implemented as a Word Document. The form was used in the DEISA DECI process and was fine tuned in DECI-7 and DECI-8 based on experience of the evaluators.

Data has to be transferred by the technical evaluator from the project's proposal submission form (another word document) to the Technical Evaluation form (word document). In the project proposal submission form no required fields are specified and some fields can be empty. This can lead to cases that the technical evaluation might be based on incomplete input.

The technical evaluation assesses the technical details of the project which can be summarized as shown below:

- what resources are required,
- what type of jobs are run,
- what is the most suitable platform for the project,
- whether the project's targets are achievable with the resources requested, and
- how much enabling/porting is required.

In the DECI-7 proposal form, applicants needed to select one platform where they wanted to run their projects. In the DECI-8 form it was requested to supply a ranking of preferred machines. This is a clear improvement, because due to both technical and accounting requirements the exec site may not be the first choice of the applicant and in fact the project may be spread over several sites. This reduces the need of home sites to contact PIs to check whether execution sites that might be proposed to them are acceptable.

The overall result of the evaluation and project feasibility can be 0, 1, or 2.

- The score 0 means that the proposed project is not feasible in the technical sense.
- The score 1 means that some support is needed by PRACE staff to port/enable the project.
- The score 2 means that no significant porting or enabling is required.

A poor score in the overall assessment in the technical review does not necessarily mean that the project will be rejected. It is the scientific review that will decide if a project will or not be rejected. Scientific reviewers take into account the technical feasibility of the project proposals and the enabling/porting effort in their decisions, though.

The peer review panel taking into consideration also the technical evaluation can recommend whether each proposal, if accepted, should be awarded the full set of resources (amount of CPU etc.) requested or whether the allocation should be scaled back.

The technical evaluation is done in two weeks, after the closing date of the corresponding DECI call. The Technical reviews are currently distributed across PRACE partners participating in DECI. Each partner reviews proposals coming from PIs of their country. This partner is named as the proposal's home site. Exceptions to this rule are countries where there are no DECI partners or if the proposal needs collaboration of more than one DECI partner. In these cases proposals are assigned based on geographical proximity of applicants, former contacts with them, and the number of proposals assigned to the sites (to balance effort).

The above-described practice leads to a rather uneven spread of workload over different partners. An alternative approach that is being considered is to split the proposals in accordance to the preferred execution platform, and form a pool of technical experts for each architecture that will do the evaluation for these applications. Due to the complexity that this process might generate it is suggested that the current approach be kept for the technical evaluation to be done by the experts of the home site.

Furthermore there has been discussion whether it would be better to conduct the technical review by execution sites, because home site staff may not be able to judge applications for architectures for which they have no experience.

An argument for leaving the technical evaluation to the home site is that the current practice has worked reasonably well for DEISA and DECI pilots.

The completed technical evaluations are being uploaded to the PRACE BSCW. These technical evaluations together with the project proposals are sent to the peer review committees. Parallel to this PRACE staff has populated the DPMDB database with the applications' technical information to be used in future.

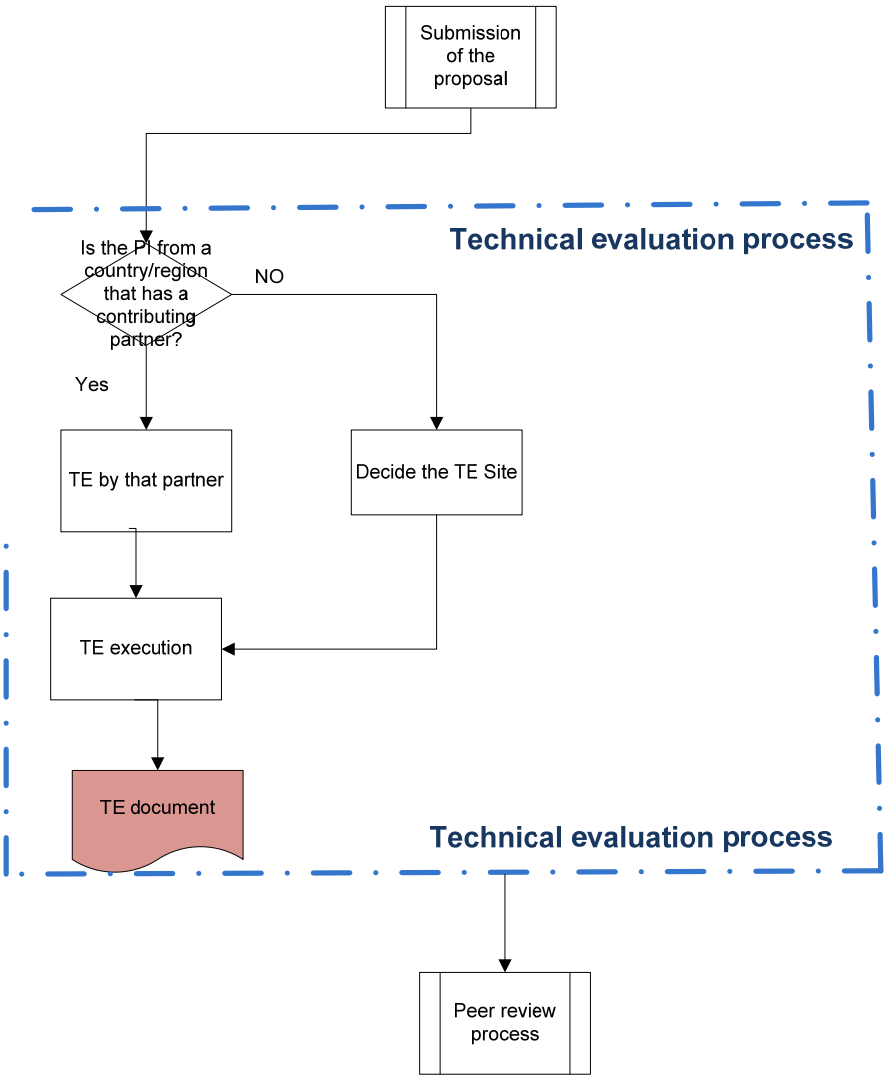

Figure 1. Technical Evaluation Process

4.5 Peer Review of Proposals

The scientific review of proposals is currently distributed across national committees in PRACE partner countries and the scientific committee of HPC-Europa2 (SUSP). Proposals coming from countries that participate in DECI are in general evaluated by national committees while the others by SUSP. Scientific panels rank proposals that are in their responsibility, but rankings of different panels are not combined. This is not necessary because resources of DECI are allocated using a *juste retour* criterion, which means that projects from PRACE partner countries obtain Tier-1 resources in proportion to how much their countries contribute with 15% of resources being reserved for external proposals (i.e. from European but non-PRACE member countries). This clearly breaks the "fairness rule" of the process, as some countries benefited independently of the quality of science.

An alternative approach would be to evaluate the proposals by a single PRACE committee, as is in the case of Tier-0 calls. With a single evaluation it would be possible to have an overall ranking of proposals with respect to all proposals submitted for the call. This would therefore mean abandoning the *juste retour* principle. The most likely consequence would be that partners might reduce their contributions for DECI, if it was perceived that the resources allocated by member countries were not available for their national scientific committees. Since the current approach works well, there seems no reason to change it.

4.6 Allocation of Resources by Partners

The call cycle and the policy of allocating resources by partners have changed over time. In DEISA, the frequency of calls initially was once per year and the execution period spanned exactly one calendar year. The minimum contribution of resources for each partner to the DECI call was 5% of the capacity of the HPC system of that partner. By the end of DEISA, overlapping calls had been introduced and the interval between calls had decreased from 12 to 9 months. The execution period of the last call in DEISA was curtailed to 8 months to fit with the end of the project funding.

In the first two DECI pilots of PRACE (DECI-7 and DECI-8), there were two calls per year. The execution period again spanned a full year, but the two execution periods of DECI-7 and DECI-8 have an overlap of 6 months.

In the policy of resource allocation the biggest change was made at the transition from DEISA to PRACE: The requirement of a minimum contribution of 5% of a partner's available resources was not enforced.

4.6.1 *Policy in DEISA DECI*

During DEISA, every partner agreed to contribute 5% of its yearly capacity to DEISA. The minimum contribution of each partner was divided over several classes:

1)15% for Virtual Communities

The remaining 85 % was divided again:

2)15% for external proposals (proposals from non-partner countries).

3) 85% for local proposals (proposals from the partner country participating in DEISA).

This division was proposed by the Application Task Force of DEISA and approved by the DEISA Executive Committee.

For DECI-1 to DECI-4 the lifetime of the projects was exactly one year and coincided with the calendar year, so matching committed resources to the DECI calls was easy. Starting with DECI-5 the lifetime of the projects was less than a year, to fit in an extra call before the end of DEISA, and execution periods didn't coincide with the calendar year anymore. Because the contributions of some sites adhere strictly to calendar years, this led to some serious arithmetic to map the committed resources to the proper DECI call.

In DEISA DECI calls the allocation of resources was based on the juste retour principle, i.e. the amount of resources that was allocated to the projects for which the partner acted as home site (the internal projects) had to be compensated by an equal contribution of resources from this partner to all the internal projects.

Every partner could volunteer to contribute extra resources, on top of the originally committed resources if the quality of their local projects was such that they wanted to award extra time to their internal projects.

4.6.2 *Policy in DECI-7 Pilot Call*

A difference with the DEISA DECI calls was that at the start of PRACE-2IP the concept of Tier-1 partner wasn't very clear yet and therefore each site was asked whether they wanted to participate in this DECI call and if they did to specify their contribution to this pilot call.

An additional difference was that in PRACE the frequency of Tier-1 calls has increased (DEISA, at the end, had one every 9 months) to twice per year. A number of partners had

specified their committed resources on a yearly basis, so they had to be split up between this DECI-7 pilot call and a future DECI-8 call, according to a formula provided by the site themselves.

As in DEISA the contribution of each partner was divided over internal and external proposals: 15% for external proposals 85% for internal proposals. In the DECI-7 and DECI-8 pilot calls there were no provisions made for community access.

As in DEISA DECI calls the allocation of resources was in kind and every partner could volunteer to contribute extra resources, on top of the originally committed resources if the quality of their local projects was such that they wanted to award extra time to their internal projects.

Unlike to the DEISA DECI calls where the requested resources from the internal proposals for each partner greatly exceeded the committed resources of that partner, in the two Tier-1 pilot calls, there were some participating sites that had more resources committed to the call than were requested by their internal projects. So they were left with some credits. In the pilot calls there was no clear policy developed on how to deal with these excess resources. Instead ad hoc decisions were made to allocate extra resources to internal projects of high scientific value that originally were cut back due to limited available resources because of the juste retour principle.

4.6.3 *Towards a Persistent Resource Commitment for PRACE DECI Tier-1 Calls*

Taking inventory who will participate in the next DECI call, prior to that call, like it was done in the DECI-7 pilot and the DECI-8 call seems to be an inefficient way to get commitments from participating sites especially because it has to be done twice a year. We should opt for longer term commitments.

We propose the following:

In principle, each partner participates in all calls. Each partner has the possibility to opt out of a particular call (for example when in a transition phase because of an upgrade or purchasing a new machine). It is the responsibility of the partner to announce his withdrawal in time before the DECI call is issued. Because, as the practice for DECI-7, DECI-8 and DECI-9 shows, there are now 2 calls per year, with a six month overlap, it will become difficult to map yearly contributions to specific DECI calls. Therefore the amount of resources that each partner contributes should be specified per call (rather than per year) with a minimum of 2.5% (or 3%?) of the yearly capacity.

In DEISA there was a close balance between committed resources, granted resources and consumed resources per site. In DECI-7 we saw for the first time that some sites contributed more to the DECI call than was granted to "their" projects and also there were accepted projects from sites that today are classified as external, but maybe in a future call want to participate as partner. To get a better balance between committed, granted and consumed resources per site, the matching of resources could be extended to include all DECI calls, i.e. accept an imbalance per call, but make sure it is compensated in a later call. This requires a meticulous bookkeeping of committed resources, granted resources and consumed resources, both for projects running on the Tier-1 system at a site and for the resources consumed by the sites "own" projects. The relevant up to date information can be found in the DPMDB. An argument against the compensations is that they add in to yearly contributions, which may mean that sites that have big deficits from previous calls must give overly large contributions in next calls, to get grants for "their" projects. Sites that have big surpluses from previous calls, in turn, may not have enough "own" projects to cut the surplus in next calls. Because it is not clear whether compensations would improve the situation or not, it is better to keep the

current practice until we have gained better insight in the consequences of extending the *juste retour* principle over multiple DECI calls.

One other important question is whether two calls per year is a good interval, or should the calls be more often or less often. Everyone agrees that the calls should not be more often than twice per year, at least not until better tools to manage the DECI process are available: the process of receiving proposals, evaluating and assigning them to execution sites takes so long that it isn't possible to shorten the DECI interval any further at the moment. A longer interval of DECI calls would have some benefits: there were more resources available per call and less work for partners. An argument against is that researchers should then wait on the average longer for calls. As long as Tier-0 calls are launched twice per year, there are no reasons to change the schedule. Having the Tier-0 and DECI calls at the same time contributed to the presentation of PRACE as a single entity.

4.7 Assignment of Resources to Accepted Projects

In the first DECI pilot of PRACE (DECI-7 call), over 90 million CPU core hours (in standardised units) were allocated from 19 machines to 35 European projects. In addition to this, large amounts of storage space and application support were allocated. Only CPU resources were balanced between sites. Other resources were considered supplementary to CPU time. Memory and storage requirements influenced the assignment process by making some sites unsuitable to certain projects.

The DECI management team reviewed accepted proposals and checked how well each machine would suit to them. The best suitable machine got a preference ranking of 1, the second best suitable a ranking of 2, and so on. Unsuitable machines got no ranking. As a result, a machine preference table was obtained. By comparing the machine preference table and resources available at machines, an assignment table was formed.

Below is a short description of the assignment process. Real projects and machines are not shown, because tables containing 19 machines and 35 projects would be too large to fit into a single page. Moreover, machine preference tables could be interpreted as overall rankings of machines, if they were published. That isn't desirable.

4.7.1 *Assignment Process*

In this short example 8 projects are assigned to 6 machines for execution. The DECI management team have reviewed the proposals and ended up in the machine preference table shown in Table 1.

Project	Allocated core-h	IBM Blue Gene/P		IBM Power6		Cray XT/XE	Cluster
		Machine A	Machine B	Machine C	Machine D	Machine E	Machine F
Proj 1	5,000,000	2	1			3	4
Proj 2	4,500,000	3	2			1	
Proj 3	4,000,000	2	1	6	5	4	3
Proj 4	3,500,000			4	3	1	2
Proj 5	3,000,000	6	5	2	1	4	3
Proj 6	2,500,000			4	3	2	1
Proj 7	2,000,000			2	1	3	4
Proj 8	1,500,000			4	3	2	1

Table 1: Machine preferences

Architecture	Machine	Available Core-h
IBM Blue Gene/P	Machine A	5,000,000
IBM Blue Gene/P	Machine B	6,500,000
IBM Power6	Machine C	2,500,000
IBM Power6	Machine D	4,000,000
Cray XT/XE	Machine E	4,500,000
Cluster	Machine F	3,500,000

Table 2: CPU resources available by machine

Even with 8 projects and 6 machines, it requires lots of brainwork to end up with a good solution. With 35 projects and 19 machines the process is many times harder. In the DECI pilot, the projects were first assigned to machine types (i.e. Blue Gene/P's, Power6's, Crays, and clusters). That helped the assignment process somewhat. This problem is a cause of not having a single priority list. All projects compete for the resources with no order. With *juste retour* it is the only option.

4.7.2 *How Should Storage and Other Resources Be Best Assigned?*

Storage requirements are a special case because they must be allocated from the same site where the project is carried out. Therefore large storage requirements may make some sites unsuitable to certain projects. To avoid wrong assignments, the technical evaluation has to be rigorous in this regard, which probably means the application/applicant needs to be more precise also.

Other resources (permanent storage, enabling effort, visualization services etc.) are easier to move to a site that has these resources available.

4.7.3 *Re-assignment of Resources*

During the DECI calls within DEISA, it turned out sometimes that resources given to projects were not suitable. If such an incident occurred, a new execution site was proposed for the project. Usually this meant that two or more projects swapped execution sites. In some cases one execution site offered to take the moving project although it meant that its total assignment grew higher than originally allocated. In DECI7 we didn't run into that problem, but this issue may arise any time in the future.

Swapping execution sites is a good alternative if the projects haven't been informed yet about their execution sites. It might also be an option later on, but then all the projects need to accept the change of execution sites. It is not easy to persuade projects to move from one execution site to another if simulations or preparations have already been started.

4.7.4 *Monitoring Resource Usage*

In DEISA DECI and the first PRACE DECI pilot both home and execution sites kept watch on projects' progression. In Tier-0's, execution sites and the AISBL do this (because Tier-0 projects have no home sites). There are some alternative ways to arrange this process: The site to keep watch on projects can be home site, exec site, enabling site, or any combination of these. The conclusion is that the current practice is good. An alternative approach could be that if a project needs enabling, the enabling site watches the project's progression. The enabling site has the closest contact with the project's PI, and is thus the best site for this task.

If the project needs only compute resources, then execution site could follow the progress, and contacts the PI or home site if needed.

There are two tools that help monitoring resource usage: DART, the Distributed Accounting Report tool, and DPMDB, the DECI Project Management DataBase. These tools are described in Secs. 4.8 and 4.9, respectively.

4.8 Distributed Accounting Report Tool (DART)

The PRACE infrastructure enables project supervisors, site administrators, and users to view computing resource usage using DART, the Distributed Accounting Report Tool. It is a Java webstart application, which gathers accounting information provided by PRACE sites in a standardised fashion. Monthly usage records are exported from DART to DPMDB, the DECI Project Management Database (Sec. 4.9).

In DECI-7 and DECI-8 DART was only implemented at the former DEISA partner sites, but not at the new sites that joined in PRACE. If we want to continue using DART as a general tool to watch the progress of DECI projects, DART should be implemented at all participating sites.

4.9 DECI Project Management DataBase (DPMDB)

To manage the DECI process within DEISA a database tool (DECI Project Management DataBase or DPMDB) was developed. In the DPMDB, personal information about the applicants and research institutions as well as technical information about the proposals and accounting information about the progress of the accepted projects was stored. This information was heavily used during the monthly videoconferences where the progress of the DECI projects was discussed.

The DPMDB was developed and maintained by RZG. DEISA staff members, who were involved in the DECI process or with the enabling work for accepted projects, all had access to the DPMDB. During the whole DECI programme, from DECI-1 to DECI-6, new features, often in a very pragmatic and ad-hoc way, were added to the DPMDB.

Because it is a very convenient tool for tracking the proposals and projects, the DPMDB was kept in the DECI-7 and DECI-8 calls. However, RZG was no longer prepared to host and maintain the DPMDB server. SARA offered to take over the hosting and maintenance of the accounting server.

The two facts that there were a lot of ad-hoc modifications to the DPMDB and that the original developer is no longer maintaining the code makes the DPMDB in its current form not very suitable for long term usage in a persistent Tier-1 environment. To make it more flexible and easy to adapt to enhancement requests the design of the database should be reconsidered.

5 Recommendations for Improving the DECI Proposal and Project Management Infrastructure

A taskforce was formed to improve DECI proposal and project management infrastructure. At a very early stage it became clear that the tool must be web-based.

The proposed web-based DECI tool that intends to support the proposal submission mechanism (Sec. 4.3), the technical evaluation process (Sec. 4.4), and the scientific evaluation process (Sec. 4.5), must provide efficient functionality to support the above processes while ensuring accountability and objectivity of the process as well as transparency and fairness to the applicants.

The following lists summarize the existing workflow and tools in the DECI process.

Existing DECI process tools

- MS Word forms are used for the submission of the proposals, the technical evaluation and scientific evaluation.
- All relevant documents are stored in BSCW.
- DECI staff and applicants communicate via e-mail.
- DPMDB (DECI Project Management Database) is a convenient tool for tracking applicant data, machine data, project data (requested resources, preferred machines, monthly usage numbers, links to documents in BSCW etc.), information from technical evaluation and scientific evaluation processes, and monitor project execution progress.

Existing DECI process workflow

- Applicant fills in and submits the application form.
- DECI staff copies (copy-paste) applicant data and project data from application form into Technical Evaluation (TE) form and assesses proposal in TE.
- Application form and TE form are sent to scientific evaluation (SE) committee, together with a template SE form partially pre filled by DECI staff and accompanied by a letter from the Project Management Office.
- SE committee fills in the SE form (pre-filled with applicant data via copy-paste from application form) and assesses the proposal.
- All forms are MS Word and/or PDF documents so far.
- DECI staff collects and uploads all forms into BSCW.
- DECI staff copies (copy-paste) applicant data and project data into DPMDB.
- An excel spread sheet is being generated to collect all the allocated projects with their resources. This sheet is used to assign projects to exec sites
- DECI staff communicates with applicants via e-mail.
- The progress of the work for obtaining certificate by users, porting their codes and further enabling is also either monitored in a spreadsheet or in a workflow tool provided by LRZ.
- Monthly resource usage numbers are copied into DPMDB (either manually by DPMDB maintainer or from DART if the exec site supports DART).

- At the end of a project report instructions are provided to users and finished reports are being gathered in BSCW.
- There is no clear procedure to track the list of scientific publications supported by DECI.

One of the possible improvements proposed is the implementation of an electronic tool (web based) to facilitate efficient, user friendly and centralized control/supervision of the proposal submission process, the technical evaluation and the scientific evaluation process.

The web-based tool will provide different views with access to relevant data to the users according to their role.

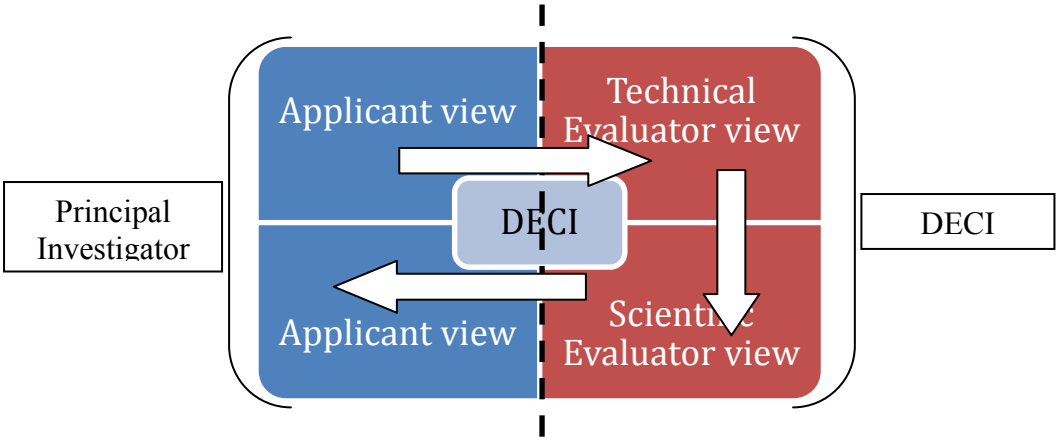

Figure 2: DECI web-based views

Content-wise, the web-based tool will mirror the fields defined in the current Word document version into web forms. It will greatly simplify the process of managing, collecting, storing and maintaining all kinds of information in a central integrated Database. That information can be either copied or linked or queried when needed into DPMDB, which will continue to be used to aid machine allocation and assignment, and monitor the project progress.

Since the application and review information will be submitted online, the submission and evaluation process would become less time consuming, scalable as well as error free due to the elimination of the current copy-paste process. The whole process would benefit from the web-based tool since many of the fields of the forms could be prefilled with information already existing from previous steps (i.e. from the application form to technical evaluation form), minimizing error occurrence. Moreover, changes in the process would be easy since the form adaptation would become an automatic process. The web-based tool would provide better communication between the applicants and the DECI staff or technical evaluators and easier document management.

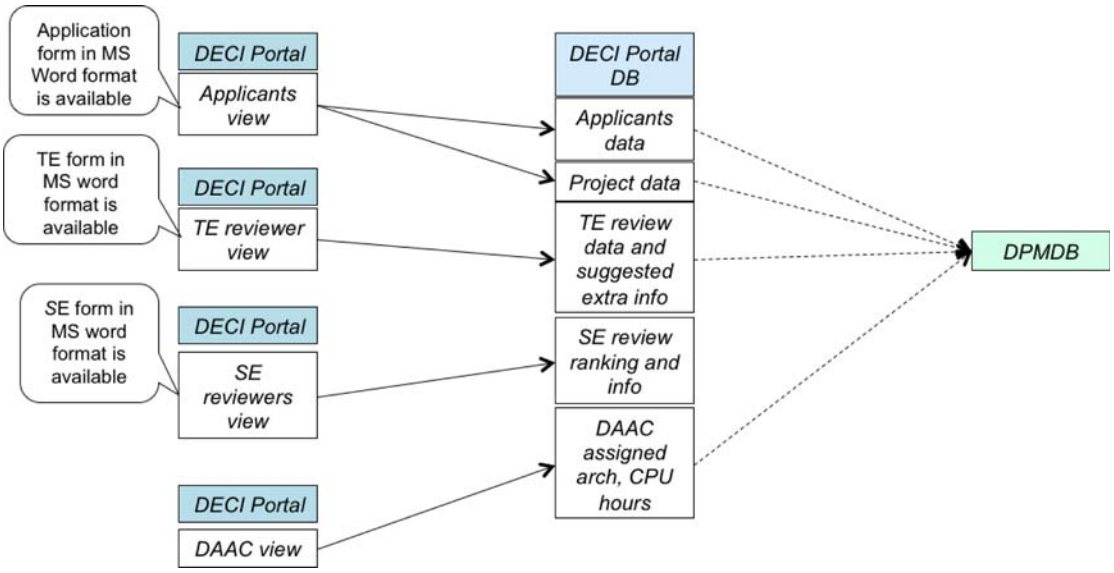

Figure 3: DECI functionality of the web based tool

The web-based DECI Tool should ideally have the following elements. The implementation of the tool can be achieved in phases by identifying some of the requirements below as essential or desirable. Essential requirements should be implemented in the first version of the tool to facilitate efficient implementation of the DECI process, while desired requirements can improve the usage of the tool and might be scheduled for implementation in subsequent versions.

- Forms Creation Flexibility:
  - Essential: Developers ability to programmatically redesign the forms contents and their integration with the internal database.
  - Desirable: Web-based ability (form design tool) to design and change the project submission and evaluation forms.
- User Workspace / Control Panel
  - Essential: Provide users with complete online control of their data (application form, user data etc.) and enable them to effectively view and browse their data (i.e. applicants can see all their applications, response letters and applications status form the portal).
- Access Right Management & Role Support
  - Essential: Assign different roles (coordinator of the process, evaluator etc.) and give access to different functionalities (i.e. evaluation assignment, evaluation process), views and data (statistical, project submission form and evaluation form) according to the different privilege level (i.e. evaluators can gain limited access to relevant proposals and TE). This would cause different log-in views for Applicants, Technical & Scientific evaluators and DAAC staff.
- Dedicated DECI Database
  - Essential: Store applicants' data, project data, TE review data, suggested extra TE info, SE data, ranking info etc. into the DECI Database.
- Administrator Functionality
  - Essential: Create and/or change user's, evaluator's, site's, countries, info.
- Reporting Support

  - Essential: Support the process of submitting a short report from the PI, after the completion of the project.
  - Desirable: Moreover the publications related to work done with DECI resources should be tracked via the proposed tool.
- Integration with the DPMDB database
  - Essential: Copy or link the relevant data from the web-based tool, when needed, into the DPMDB (i.e. project name, home site, technical requirements such as CPU type, number of jobs, memory, simulation codes etc.).
  - Desirable: Copy summary of projects' resource usage from DPMDB to the web-based tool, so that PIs can view accounting information related to their projects without learning a new tool (DART).
- Export Support
  - Desirable: Create and export documents and information that should feed other systems or processes (i.e. automatic generation and export of PDF's for mailing at any point in time)
- Logging
  - Desirable: Keep extensive logs regarding all changes made by the users in the tool.
- Communication Support
  - Desirable: Provide different communication tools (via email, via user workspace etc.) between the users who have to communicate according to the existing workflow (i.e. technical evaluator and principal investigator).
- Workflow Support
  - Desirable: Design and run workflows between the Coordinators of the Evaluation Process, the evaluation sites and the evaluators. The web-based DECI tool could support rule creations that would be associated with conditions and actions (i.e. time reminders or enforcement – establish deadlines for submission of evaluation, email reminders to reviewers, alerts to the evaluators of completed, pending or overdue reviews).
- Statistical Report Support
  - Desirable: Create statistics reports of the DECI process (i.e. number of technical evaluations per site, number of scientific evaluations per evaluator).

6 Issues of Cross-national Programme for Resource Exchange for Community Access

6.1 General Issues

As the Community Access for Tier-1 has not been decided yet, there are no practical experiences of it and the issues are related mostly to the preparations of the access.

The more political issues of Scientific Communities could be discussed within the framework of the envisioned PRACE Tier-1 Optional Programme of the PRACE RI.

6.2 Access Time for Communities

From the point of view of most applicants, the period of access for Tier-1 Community Access should be similar to PRACE Tier-0 multi-year access, and the process should be harmonized as much as possible.

It would be easier for the communities if the calls could be synchronized with the Tier-0 multi-year programme calls, and if they would be equally long in time. It would be beneficial if the applicants could combine Tier-0 and Tier-1 resources in their applications.

6.3 Resources for Community Access

The Tier-1 resources for the virtual communities could be taken from the same pool as the DECI projects. It might not be necessary to decide at the start of the allocation period about all future CPU and other resources to be committed to the communities as long as a commitment to provide them with appropriate access was made for a longer time. The decision on the actual amount of resources could be made, e.g. in conjunction with the allocation of resources for DECI projects. A share of all available resources, like in DEISA, is also possible.

6.4 Review Process for Community Access

The reviewing process should be approved by the participating sites. The evaluation could be conducted in the same way as with external DECI proposals:

- technical review by participating sites (the site nearest to the PI, either geographically or organisationally), and
- scientific peer review by the scientific committee of HPC-Europa2 (SUSP).

There should also be a light reporting and review cycle every year for the communities.
The possibility of Tier-1 scientific community access should be taken in to account when the peer review tools for DECI are designed and implemented.

7 Issues of Cross-national Programme for Resource Exchange for Industry Access

7.1 Overview

As was remarked in Sec. 2.3, PRACE regards the provision of HPC services to industry a very important part of its mission. Since January 2012, the PRACE AISBL Council allowed industrial users to apply to PRACE Tier-0 calls through Open R&D projects with the following conditions:

- Industrial users will apply on the same calls as users from academia (Preparatory Access and Regular calls);
- PRACE AISBL resources are accessible for companies seated in Europe or having a significant R&D activity in Europe;
- Eligible companies commit to publish results obtained by the use of PRACE resources at the end of their grant period;
- Companies may apply alone or paired with partners from academia. In the first case the amount of resources asked for by a company will be limited to 5% of a single system;
- For regular calls, the only selection criterion will be scientific excellence while for preparatory access calls proposal will be evaluated only technically.

Based on those conditions, access to PRACE Tier-0 resources will be free of charge for selected eligible companies.

7.2 Integration in PRACE Tier-1 resources allocation process

Without pretending to be exhaustive, a survey conducted by WP5 of PRACE-1IP, covering 34 industrial representatives, shows that the European HPC industrial needs for efficient runs typically require 256 cores. Hence one can assume that the average industrial users needs will be met to a higher degree by Tier-1 systems than Tier-0 ones.

In addition, almost 85% of survey representatives answered that the cost of the licenses for commercial applications or code scaling/development for in-house or open source applications prevents them from using HPC or extending the use of HPC. Combining this information with data that 97 % of the industrial companies that employ HPC consider it indispensable for their ability to innovate, compete, and survive [4], it seems clear that a future PRACE role in promoting and strengthening European companies' competitiveness by awarding them access to its computational resources and in general to HPC competence both on Tier-0 and Tier-1 level has been identified.

This assumption towards PRACE industrial engagement is based on the identified requirements of the industrial European users in the area of HPC by nature [6]:

- Evangelisation[3], information, and training;
- On-demand access to HPC resources;
- Expertise in scientific software;
- Code enabling;

[3] Targeted marketing by a HPC centre

- Co-design of industrial applications;
- An open R&D programme;
- Pre-competitive R&D[4];
- Partnership with HPC research centres;
- Technology transfer and other high-value knowledge services;
- Commercial activities.

In order to provide industrial access at Tier-1 level, there is a need by PRACE to modify its operational procedures in order to support not only capability-based jobs (a single run on a high number of cores) but also those that are capacity-based (multiple executions on a low number of cores). This latter category will benefit industrial users dealing with coupled multiphysics codes or managing uncertainties.

7.2.1 *DECI Industry Use*

Following the assessment of the needs and expectations of European industrial users it seems that PRACE Tier-1 level of resources and services is inevitable to engage and to propose high value services for industry. Of course, this may be regarded as just a step for future usage of Tier-0 systems.

Based on the project access to Tier-1 resources, via the DECI described above (Sec. 4), a single Open R&D project cross-national access may be provided for industrial users to European Tier-1 resources. It may start with the same rules as obtaining access to Tier-0 systems. Companies may apply on their own or paired with partners from academia.

As with academic projects, each successful project will be awarded cross-national access to European Tier-1 resources for a period of 12 months, beginning from an initially-defined date. Open R&D access will enable European industrial researchers to obtain access to a fraction of the most powerful national (Tier-1) computing resources in Europe, regardless of their country of origin or work and thus 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.

Similarly to other DECI projects, all industrial open R&D projects will be technically and scientifically peer reviewed. The DECI proposals submitted by a PI affiliated to an industrial organization located in one of the Tier-1 countries will be peer-reviewed by their own national peer-review system, as for the scientific projects.

Additionally, PRACE needs to investigate the other access modes such as pre-competitive projects or those offering small-scale services5 for commercial activities or studies. The adoption of a type of industry access other than Open R&D, like those for pre-competitive projects or those offering small-scale services for commercial activities or studies would have to be taken with caution within PRACE, since questions like distortion of the market rules as set by national state aid laws access may rise. If an industry need of this kind is identified, this probably should be addressed first in PRACE AISBL and the corresponding legal advice should be sought for setting a PRACE policy in this regard.

In the future, besides offering access to computing resources, giving applications-enabling assistance from experts at the leading European HPC centres should be investigated. This will be tested already in the PRACE-2IP project by providing an offer to enable projects to be run on the most appropriate Tier-1 platforms in PRACE.

[4] Early stages of the development of a commercial product, during which competitors collaborate
[5] E.g. an application portal to be used by a commercial company

The DECI proposals submitted by PI affiliated to an industrial organisation located in countries not acting as Tier-1 will be peer-reviewed by the HPC-Europa2 SUSP Panel as above. Such external proposals can apply for porting and optimising industrial applications only for open research and will not be permitted for pre-competitive projects, or those offering small-scale services for commercial activities or studies.

In a similar way as for Tier-0, for Tier-1 calls proposals from industry are eligible, as long as the project leader is a senior researcher employed in an industrial organisation seated in a European Union country or has a substantial R&D activity in Europe. Under PRACE rules for Tier-0, the employment contract of the project leader with the industry organisation must be valid until at least three months after the end of the allocation period.

PRACE HPC centres may have further restrictions on who is eligible to use the machines including the admissibility of pre-competitive projects or the offering of small-scale services for commercial activities or studies, as well as legal restrictions (for example due to export rules or security measures). It is the responsibility of the applicant to ensure that the company is eligible to use the system.

8 Issues with Recording and Monitoring Usage

There are two tools that help monitoring resource usage: DART, the Distributed Accounting Report tool (Sec. 4.8), and DPMDB, the DECI Project Management DataBase (Sec. 4.9). To access accounting data with DART a user must have a valid PRACE account. The data of DPMDB is viewed by the DECI management team only. Monthly usage records are exported from DART to DPMDB.

There are two issues related to this procedure:

1) PIs do not always have PRACE accounts, so they don't have access to DART either. Therefore site administrators need to send accounting data to the PIs by email. It requires some manual work, because DART doesn't show monthly records in a single view. The data on the DPMDB is in a more convenient format, but it is not updated as frequently as some might like.
2) Accounting information is composed in a different way by DART and sites' own accounting report tools. Therefore the usage records may differ. The differences are usually small, but it takes time to explain to PIs why they see different numbers with different tools, and which tool they should trust.

An easy solution to the first problem would be to present accounting information in the DECI tool, which is used by PIs to submit proposals and to follow review process. In that way, all the most important information needed by PIs would be in the same place.

9 Issues with Over or Under Utilisation of Resources

In the last finished DECI (DECI-6 of DEISA), 91.6 million standardised CPU-core hours were granted and 94.4 million were consumed. This means 3% overdraft, which is totally fine. However, if we look deeper in the statistics we see that most projects either over or under utilise their project's budget:

- Of 56 projects, 32 used fewer resources than were granted. The average usage of those projects was 57%.
- The other 24 projects used more resources than were granted to them. The average usage of those projects was 140%.

Similar phenomenon is seen in the site statistics:

- Of 12 sites, 9 were underutilised. The average usage at those sites was 73%.
- The other 3 sites were overdraft. The average usage at those sites was 133%.

Underlying causes for the over and under utilisation are:

- Some sites do not have a mechanism to limit resource usage, or the mechanism is not compliant with PRACE procedures, e.g. limit usage by a UNIX group rather than PRACE project. (overutilisation)
- The computing time requirement of the research project was over or under estimated (over/underutilisation)
- Work management in research groups (underutilisation)
- Technical problems (underutilisation)

Why this is a problem?

- Sites are not happy because the *juste retour* criterion isn't met.
- Projects are not happy if underutilisation takes place because of external reasons.

However, a degree of flexibility which enables resources to be transferred between projects, according to their needs, makes best use of the available resources and maximises the overall science output. Many sites are happy to commit to an overutilisation in order to support successful projects, by prior agreement.

All sites have local policies employed how to deal with under- and overutilization. Currently, DECI (as well as the Tier-0 access) relies on these local policies, and does not define its own policies.

10 Conclusions

It is already clear that the role and purpose of Tier-1 in the PRACE infrastructure is evolving. There are three main drivers for change:

1) The number of Tier-0 systems is growing, as is the range of architectures available at Tier-0. Increasingly, Tier-0 can offer a Petaflop version of each of the Tier-1 architectures. This means that there is a much clearer migration path from Tier-1 to Tier-0 than was apparent when PRACE-1IP started.
2) The number of Tier-1 systems available in PRACE has continued to grow with the vast majority of the PRACE partners now participating in resource exchange. The majority of the new partners have small systems, so the average size of a Tier-1 system has dropped. The total amount of CPU resources available however, although more in total than was available in DEISA2, is split over many more systems.
3) As Tier-1 systems come to the end of their lifecycle, they are typically replaced by much larger systems. The larger countries tend to have bigger budgets for HPC and are able to replace or upgrade systems more often, meaning that the divergence in size between the smallest and the largest Tier-1 systems in PRACE is also growing.

The challenge for PRACE is therefore to come up with a model for resource exchange or resource pooling which takes account of the changing circumstances. Where the pool of Tier-1 resources comprises a number of similar-sized machines of different architectures (as was the case in DEISA and as is the case with Tier-0), the benefits of resource sharing to both resource providers and resource users are easy to explain. The attractions of having access to alternative architectures or machine configurations of similar sized machines for scientific codes with different characteristics are obvious. However, within Tier-1 in PRACE, at the moment, some of the resources are smaller than the traditional DECI user communities are used to. This has prompted a consideration of how PRACE can best optimise the use of the resources at its disposal. Large systems and novel systems (e.g. GPU resources, high memory clusters etc.) are of particular interest to the DECI community. Medium sized clusters (i.e. machines similar to but of an order of magnitude larger than typical University departmental facilities) are also popular with DECI researchers making their first steps from local to European facilities.

As the scope of Tier-1 activities grows, however, it will become easier to utilise the many smaller Tier-1 machines in an optimal way. There is an opportunity to use these machines to support Tier-1 activities new to PRACE – in particular, to foster industrial use of PRACE facilities, since many small to medium-sized enterprises, and even larger companies that could benefit from HPC do not have codes that can currently scale to thousands of processors. Given the lower starting point and levels of expertise of the vast majority of potential industrial users of HPC, it is important to provide them with ramps not only to Tier-0, but also to the larger Tier-1 facilities. Smaller Tier-1 facilities could play an important role in this. Additional potential areas of interest where Tier-1 facilities could be very useful are:

- in training, where we expect PRACE's activities to grow considerably over the next few years;
- in providing development resource for science communities; and
- in supporting multi-scale applications.

These are all areas which warrant further exploration.

We recommend that PRACE considers the development of a Tier-1 strategy, focussing both on the role of a Tier-1 pool of contributed resource in supporting single-project access, science communities, industrial access, training and code development and testing. This strategy should cover not only the role of Tier-1 as an entity in itself, but also its role in the wider HPC ecosystem, with particular reference to its importance in supporting migration from Tier-1 to Tier-0. We suggest that this strategy and the subsequent management of activities could best be organised and supported within the framework of an Optional Programme (see Sec. A).

Resource exchange or resource pooling lies at the heart of a successful Tier-1 Optional Programme – without contributed resources in sufficient quantity or of sufficient quality, the Optional Programme would be unable to deliver on its objectives. Without a fair and transparent resource exchange mechanism, which can provide an objective evaluation of an increasingly wide and divergent set of resources provided by the PRACE partners, they will be increasingly reluctant to contribute these resources to PRACE.

The conclusion of this document is that the processes of DECI, which were migrated nearly "as such", are fit for their current purpose. However, the pilot identified a number of areas in which these processes (and the tools which support and sustain their efficient operation) could be improved, with two particular aims in mind. Firstly, the creation of processes better able to cope with a divergence of resources and an increase in the number of partners and secondly, where appropriate, to provide better synergy with Tier-0 for the benefit of both providers and users of PRACE resources.

For example, at present, Tier-0 and Tier-1 calls are synchronised but separate. In Tier-0 calls, an online submission mechanism is used, while in Tier-1, proposals are sent as email attachments. A recommendation is that Tier-0 and Tier-1 calls are harmonised as far as possible, so that applicants do not need to specify which Tier they direct their proposals to, unless they wish to. The submission mechanism of Tier-0 is more professional-looking and easier to use than that of Tier-1, and therefore the former would be good to use in both calls.

Scientific evaluations of Tier-1 proposals are distributed across national committees in PRACE partner countries, with few exceptions. The most important exception is external proposals – proposals from countries that do not contribute Tier-1 resources for PRACE. External proposals are evaluated by an independent review panel: HPC-Europa2 SUSP. Resources are allocated to projects using a juste retour criterion. That means that projects from PRACE partner countries get Tier-1 resources in proportion to how much their countries contribute. 15% of resources is reserved for external proposals. The current practice of how proposals are evaluated and how resources are allocated works well for DECI, and no changes are suggested. However, change in the provision of peer-review will be needed when the HPC-Europa2 project comes to an end in December 2012, so alternative arrangement for Tier-1 peer review should be explored as a matter of some urgency. In addition, allocation of Tier-1 resources may need to be re-considered when making arrangements for industrial access.

Resources are granted to projects in "PRACE standard core hours". The actual amount of core hours that a project can use depends on the supercomputer architecture the project is assigned to. Conversion factors between architectures are used. Most of the conversion factors were determined during DEISA by running a number of applications from the DEISA Benchmark Suite. An alternative approach would be to match conversion factors with cost prices of a core hour on architectures. Both approaches, benchmark- and cost-based, have their own strengths and weaknesses, and there is no single answer as to which is better. But in any case, conversion factors are essential for resource exchange.

Assignment of resources to projects is done by PRACE partners. It is a complex and tedious task, where some kind of visual tool would be useful. Although a committee-based assignment process has its disadvantages, it is a common practice, and no good alternatives to it were found.

The document hopefully contains useful information for the council of PRACE AISBL, and all PRACE-2IP personnel involved with DECI. The activities begun in task 4.3 of PRACE 1-IP will be continued in WP2 of PRACE 2-IP where the implementation of the recommendations, if accepted by the AISBL, will be carried out.

A. Annex

REGULATION NR. 2 [Partnership For Advanced Computing In Europe] OPTIONAL PROGRAMMES

BACKGROUND:

A. Words and expressions used in this regulation in capital letters have the same meaning as the definitions set out in article 1 of the statutes of the Association ("Statutes"); in case of contradiction between this Regulation and the Statutes, the Statutes shall prevail;

B. This Regulation shall be approved in a Council's meeting of the Association.

Article 1

(Activities of the Association)

1. The Association shall carry out mandatory activities - which are necessary for the achievement of its purposes - as described in article 3, paragraph 2 of the Statutes, and may carry out optional activities which may not be detrimental to the mandatory activities of the Association.
2. In accordance with the provisions set forth herein, the Association shall ensure the execution of the following particular Optional Programmes:

Article 2

(Proposal and Participation of Members in an Optional Programme)

1. If a proposal for the carrying out of an optional programme covered by article 1, paragraph 2 above is made, the Chairman of the Board of the Council shall communicate it to all Members for analysis.

2. Once the Council has, in accordance with article 14, paragraph 3, sub-paragraph d) of the Statutes, approved the carrying out of an optional programme within the framework of the Association, the Members that intend to take part in such an optional

programme shall, within a period of three months, formally declare that they are interested in participating in the same by means of a letter addressed to the Chairman of the Council of the Association.

3. The participating Members will compose a specific committee ("the Optional Programme Committee") which shall be entitled to propose to the Council all the terms and conditions under which an optional programme is to be performed.

4. The participating Members shall draw up a declaration ("Declaration") which, subject to article 4, paragraph 4 below, shall set out their undertaking in respect of:

   a) the phases of the programme;

   b) the conditions under which it is to be carried out, including the timing, the indicative financial envelope and sub-envelopes relating to phases of the programme and any other provisions for its management and execution;

   c) the scale of contributions as per article 4, paragraph 3 of this Regulation;

   d) the duration and amount of the first binding financial commitment.

5. The Declaration shall be provided to the Council for information, together with a draft on the implementation rules of the optional programme ("Implementation Rules") to be approved by the Council.

6. If a participating Member is unable to accept the provisions set out in the Declaration and in the Implementation Rules as approved by the Council within the timeframe established in the Declaration, it shall cease to be a participating Member of such optional programme.

7. Non participating Members may subsequently become participating Members by accepting the provisions included in the Declaration and in the Implementation Rules in accordance with conditions to be determined with the participating Members.

Article 3

(Execution of Optional Programmes)

1. Without prejudice to article 14, paragraph 3 of the Statutes, decisions of the participating Members regarding optional programmes shall be taken in accordance with this Regulation and the Implementation Rules.

2. The optional programmes shall be executed in accordance with the provisions of this Regulation, the Implementation Rules and the Statutes.

Article 4

(Budget, Contribution of Members and Costs)

1. A balanced budget shall be prepared by the Optional Programme Committee.

2. Each Member shall contribute to the costs of each optional programme, if it has formally declared itself interested in participating as per paragraph 2 of article 2 above.

3. The participating Members shall unanimously decide on the scale of contributions to a given programme (article 2, paragraph 4, sub-paragraph c) above) and its governance rules.

4. If a programme includes a project definition phase, the participating Members shall, at the end of such phase, reassess the cost of the programme. If the reassessment shows that there is a cost overrun greater than 20% of the indicative financial envelope referred to in article 2, paragraph 4, sub-paragraph b) above, any participating Member may withdraw from the programme. The participating Members that wish to continue with the programme, notwithstanding the cost overrun, shall determine the arrangements for such continuation. Once determined, they shall report to the Council, which shall take any and such measures as may thereby be required.

5. The Council shall establish a procedure enabling the financial envelope or sub-envelopes to be revised in the event of price level variations.

Article 5

(Termination of membership)

Termination of the membership of a Member as per paragraphs 1 or 2 of article 10 of the Statutes entails the withdrawal of that Member from all the optional programmes in which it participates. This is without prejudice to paragraphs 3 to 5 of article 10 of the Statutes.

Article 6

(End of an Optional Programme)

1. The Optional Programme Committee may decide to discontinue a programme.

2. The Chairman of the Council of the Association shall notify the participating Members of the completion of a programme in accordance with the Implementation Rules., which shall cease to be in force upon receipt of such notification.