



**SEVENTH FRAMEWORK PROGRAMME  
Research Infrastructures**

**INFRA-2011-2.3.5 – Second Implementation Phase of the European High  
Performance Computing (HPC) service PRACE**



**PRACE-2IP**

**PRACE Second Implementation Phase Project**

**Grant Agreement Number: RI-283493**

**D2.3  
Second Annual Report of WP2**

***Final***

Version: 1.1  
Author(s): Juha Fagerholm, CSC, and Alison Kennedy, EPCC  
Date: 22.8.2013

## Project and Deliverable Information Sheet

PRACE Project	<b>Project Ref. №:</b> RI-283493	
	<b>Project Title:</b> PRACE Second Implementation Phase Project	
	<b>Project Web Site:</b> <a href="http://www.prace-project.eu">http://www.prace-project.eu</a>	
	<b>Deliverable ID:</b> < D2.3 >	
	<b>Deliverable Nature:</b> <DOC_TYPE: Report >	
	<b>Deliverable Level:</b> PU *	<b>Contractual Date of Delivery:</b> 31 / August / 2013
		<b>Actual Date of Delivery:</b> 31 / August / 2013
<b>EC Project Officer:</b> Leonardo Flores Añover		

\* - 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	<b>Title:</b> Second Annual Report of WP2	
	<b>ID:</b> D2.3	
	<b>Version:</b> <1.1 >	<b>Status:</b> <i>Final</i>
	<b>Available at:</b> <a href="http://www.prace-project.eu">http://www.prace-project.eu</a>	
	<b>Software Tool:</b> Microsoft Word 2007	
	<b>File(s):</b> D2.3.docx	
Authorship	<b>Written by:</b>	Juha Fagerholm, CSC, and Alison Kennedy, EPCC
	<b>Contributors:</b>	Michael Browne, JCHEC; Xu Guo, EPCC; Chris Johnson, EPCC; George Karagiannopoulos, GRNET; Ioannis Liabotis, GRNET; Petri Nikunen, CSC; Marianthi Polydorou, GRNET; Andreas Polyraakis, GRNET; Georgi Prangov, NCSA; Wim Rijks, SURFsara
	<b>Reviewed by:</b>	Stephane Requena, Genci; Veronica Teodor, FZJ
	<b>Approved by:</b>	MB/TB

## Document Status Sheet

Version	Date	Status	Comments
0.1	19/July/2013	Draft	2.1 (partly), 2.3, 3
0.2	25/July/2013	Draft	1, 2.0, 10.1 added
0.3	29/July/2013	Draft	2.1.3, 2.2,4,5,6,7,8 added
0.4	5/August/2013	Draft	Exec summary added
0.5	5/August/2013	Draft	Additions from Alison

1.0	8/August/2013	Final version	For internal review
1.1	22/August/2013	Final version	Final

## Document Keywords

<b>Keywords:</b>	PRACE, HPC, Research Infrastructure
------------------	-------------------------------------

### Disclaimer

This deliverable has been prepared by the responsible Work Package of the Project in accordance with the Consortium Agreement and the Grant Agreement n° RI-283493. It solely reflects the opinion of the parties to such agreements on a collective basis in the context of the Project and to the extent foreseen in such agreements. Please note that even though all participants to the Project are members of PRACE AISBL, this deliverable has not been approved by the Council of PRACE AISBL and therefore does not emanate from it nor should it be considered to reflect PRACE AISBL's individual opinion.

### Copyright notices

© 2013 PRACE Consortium Partners. All rights reserved. This document is a project document of the PRACE project. All contents are reserved by default and may not be disclosed to third parties without the written consent of the PRACE partners, except as mandated by the European Commission contract RI-283493 for reviewing and dissemination purposes.

All trademarks and other rights on third party products mentioned in this document are acknowledged as own by the respective holders.

## Table of Contents

Project and Deliverable Information Sheet .....	i
Document Control Sheet.....	i
Document Status Sheet .....	i
Document Keywords .....	ii
Table of Contents .....	iii
List of Figures .....	iv
List of Tables.....	iv
References and Applicable Documents .....	iv
List of Acronyms and Abbreviations.....	v
Executive Summary .....	1
<b>1 Introduction .....</b>	<b>1</b>
<b>2 Work Undertaken in Project Months 13–24.....</b>	<b>2</b>
<b>2.1 Framework for Resource Exchange.....</b>	<b>3</b>
2.1.1 <i>Develop an administrative framework for resource exchange between Tier-1 sites to support DECI calls (T2.1.1)</i> .....	4
2.1.2 <i>Analyze options for a possible exchange between Tier-0 and Tier-1 sites (T.2.1.2)</i> .....	6
2.1.3 <i>Refine the pilot practices developed by PRACE-IIP towards production level practices (T2.1.3)</i> 8	8
2.1.4 <i>Create contracts, processes and exchange policies (T2.1.4)</i> .....	8
2.1.5 <i>Study new usage models for Tier-0 and Tier-1 resources (T2.1.5)</i> .....	9
2.1.6 <i>Investigate flexible resource exchange policies for load balancing between sites and over time (T2.1.6)</i> .....	10
<b>2.2 DECI process .....</b>	<b>15</b>
2.2.1 <i>DECI-7</i> .....	16
2.2.2 <i>DECI-8</i> .....	16
2.2.3 <i>DECI-9</i> .....	16
2.2.4 <i>DECI-10</i> .....	16
2.2.5 <i>DECI-11</i> .....	17
2.2.6 <i>DPM Database</i> .....	17
2.2.7 <i>Scientific evaluation forms</i> .....	17
<b>2.3 Community Access .....</b>	<b>17</b>
<b>3 Collaboration with PRACE-2IP WP7 Task 7.2.....</b>	<b>19</b>
<b>4 Collaboration with PRACE-3IP WP2 .....</b>	<b>19</b>
<b>5 Overview of DECI-9 Projects.....</b>	<b>20</b>
5.1.1 <i>Projects by application area</i> .....	20
5.1.2 <i>Projects by country of origin</i> .....	20
5.1.3 <i>Resources Allocated</i> .....	21
<b>6 Overview of DECI-10 Proposals and Projects.....</b>	<b>22</b>
6.1.1 <i>Projects by application area</i> .....	22
6.1.2 <i>Projects by country of origin</i> .....	23
6.1.3 <i>Resources Allocated</i> .....	24
<b>7 Overview of DECI-11 .....</b>	<b>25</b>
7.1.1 <i>Overview of DECI proposals for all calls</i> .....	25
<b>8 Analysis of Tier-1 Resources .....</b>	<b>25</b>

9	Annex.....	27
9.1	Tier-1 preparatory access form requirements .....	27

### List of Figures

Figure 1	Timeline of DECI calls showing those dealt with under PRACE-2IP and PRACE-3IP .....	15
Figure 2	DECI-9 accepted projects shown by scientific discipline .....	20
Figure 3	DECI-10 projects shown by percentage accepted by scientific discipline.....	23
Figure 4	Number of proposals received for all DECI calls so far under both DEISA and PRACE.....	25
Figure 5	Resources contributed per DECI call.....	26

### List of Tables

Table 1	Proposals received for DECI-9 together with proposals accepted .....	20
Table 2	DECI-9 Proposals and Projects by country of origin .....	21
Table 3	DECI-9 Projects by home site showing execution site and resources awarded .....	22
Table 4	Proposals received for DECI-10 together with proposals accepted .....	22
Table 5	DECI-10 Proposals and Projects by country of origin .....	23
Table 6	DECI-10 Projects by home site showing execution site and resources awarded .....	24
Table 7	Resources contributed by partner in thousands of standard hours .....	25

### References and Applicable Documents

- [1] <http://www.prace-project.eu>
- [2] PRACE Digest 1/2013, [http://www.prace-ri.eu/IMG/pdf/prace\\_digest\\_2013.pdf](http://www.prace-ri.eu/IMG/pdf/prace_digest_2013.pdf)
- [3] Deliverable D2.1 “Migration from DEISA2 to PRACE-2IP”
- [4] Deliverable D2.2 “First Annual Report of WP2”
- [5] PRACE-1IP Deliverable D4.3.2 “Cross-National Programme for Tier-1 Access Pilots”
- [6] George Karagiannopoulos, Ioannis Liabotis, and Wim Rijks “An Administrative Framework for Tier-1 Resource Exchange within Europe”, PRACE White paper
- [7] COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT, THE COUNCIL, THE EUROPEAN ECONOMIC AND SOCIAL COMMITTEE AND THE COMMITTEE OF THE REGIONS: High-Performance Computing: Europe's place in a Global Race – Brussels, 15.2.2012 COM(2012) 45 final
- [8] PRACE Scientific Case for HPC in Europe 2012–2020, <http://www.prace-ri.eu/PRACE-The-Scientific-Case-for-HPC>

All deliverables are available at <http://www.prace-ri.eu/Public-Deliverables>

## List of Acronyms and Abbreviations

AISBL	Association Internationale Sans But Lucratif (International non-profit organisation); PRACE AISBL
BSC	Barcelona Supercomputing Center (Spain)
BSCW	Basic Support for Cooperative Work; a web based system that offers shared workspaces
CINECA	Consorzio Interuniversitario; the largest Italian computing centre
CINES	Centre Informatique National de l'Enseignement Supérieur (The National Computer Centre for Higher Education, represented in PRACE by GENCI, France)
CPU	Central Processing Unit
CSC	Finnish IT Centre for Science (Finland)
CSCS	Centro Svizzero di Calcolo Scientifico (The Swiss National Supercomputing Centre, represented in PRACE by ETHZ, Switzerland)
Cyfronet	A Polish computing centre
DAAC	DECI Access and Allocations Committee
DART	DECI Accounting Report Tool; a Java web-start application to retrieve accounting data from computing facilities
DECI	Distributed European Computing Initiative; a scheme through which European researchers 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
DoW	Description of Work (PRACE-2IP)
DPMDB	DECI Project Management Database; a web-based application to view and edit details of DECI proposals and projects
EC	European Commission
EPCC	Edinburg Parallel Computing Centre (represented in PRACE by EPSRC, United Kingdom)
EPSRC	The Engineering and Physical Sciences Research Council (United Kingdom)
ETHZ	Eidgenössische Technische Hochschule Zürich (ETH Zurich, Switzerland)
FZJ	Forschungszentrum Jülich (Jülich Supercomputing Centre, Germany)
GCS	Gauss Centre for Supercomputing (Germany)
GENCI	Grand Equipement National de Calcul Intensif; the French HPC agency
GHz	Gigahertz; 10 <sup>9</sup> clock cycles per second
GFlop/s	10 <sup>9</sup> floating point operations per second (usually in 64-bit, i.e. double precision)
GPGPU	General Purpose GPU
GPU	Graphic Processing Unit
GRNet	Greek Research and Technology Network
HLRS	Hochleistungsrechenzentrum Stuttgart (High Performance Computing Center Stuttgart, represented in PRACE by GCS, Germany)
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with supercomputing
HPC-Europa	An EU-funded programme by which scientists make short research visits to collaborate with a research department working in a similar field
ICHEC	Irish Centre for High-End Computing (represented in PRACE by NUI Galway)

ICM	Interdyscyplinarne Centrum Modelowania Matematycznego i Komputerowego (The Interdisciplinary Centre for Mathematical and Computational Modelling, Poland)
IDRIS	Institut du Développement et des Ressources en Informatique Scientifique (Institute for Development and Resources in Intensive Scientific computing, represented in PRACE by GENCI, France)
KTH	Kungliga Tekniska Högskolan (Royal Institute of Technology, represented in PRACE by SNIC, Sweden)
LINPACK	Software library for Linear Algebra; LINPACK benchmarks are a measure of a HPC system's floating point computing power
LRZ	Leibniz-Rechenzentrum (Leibniz Supercomputing Centre, represented in PRACE by GCS, Germany)
MB	Management Board (PRACE-2IP)
MS	Microsoft; an international software company; or Milestone (PRACE-2IP)
NCF	Netherlands Computing Facilities (the Netherlands)
NCSA	National Centre for Supercomputing Applications (Bulgaria)
NUI	National University of Ireland
OP	Optional Programme, a programme in which members' participation is optional, defined in PRACE AISBL Statutes
PDC	Center for High Performance Computing, at KTH (represented in PRACE by SNIC, Sweden)
PDF	Portable Document Format
PI	Principal Investigator
PMO	Project Management Office (PRACE-2IP)
PPR	Project Proposal and Reporting
PRACE	Partnership for Advanced Computing in Europe
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)
RZG	Rechenzentrum Garching (Garching Computing Centre, of the Max Planck Society, represented in PRACE by GCS, Germany)
SARA	Stichting Academisch Rekencentrum Amsterdam; a Dutch computing centre (represented earlier in PRACE by NCF, Netherlands)
SE	Scientific Evaluation of a DECI proposal
SHAPE	SME HPC Adoption Programme in Europe
SIDB	System Information Database
SIGMA	The Norwegian Metacenter for Computational Science
SME	Small and medium-sized enterprises
SNIC	Swedish National Infrastructure for Computing
STFC	Science and Technology Facilities Council (represented in PRACE by EPSRC, United Kingdom)
SSC	Scientific Steering Committee (PRACE)
std-hour	Standard CPU core hour; Corresponds roughly one hour on an IBM Power4+ 1.7 GHz processor
SURFsara	Dutch national High Performance Computing & e-Science Support Center
SUSP	Scientific Users' Selection Panel (HPC-Europa)
TASK	Trojmiasto Academic Computer Network (Poland)

TB	Technical Board
TE	Technical Evaluation of a DECI proposal
TFlop/s	1012 floating point operations per second (usually in 64-bit, i.e. double precision)
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
UHeM	Istanbul Technical University National Center for High Performance Computing (=UYBHM)
UiO	Universitetet i Oslo (represented in PRACE by SIGMA, Norway)
UYBHM	Ulusal Yüksek Başarılı Hesaplama Merkezi (The National Center for High Performance Computing, Turkey)
VSB-TUO	Vysoká škola báňská – Technická univerzita Ostrava (Technical University of Ostrava, Czech Republic)
WCNS	Wroclaw Centre for Networking and Supercomputing (Poland)
WP	Work Package



## Executive Summary

This document, the “Second Annual Report of WP2”, of the PRACE Second Implementation Phase Project (PRACE-2IP) reports on the work undertaken in the final 12 month period from 1 September 2012 to 31 August 2013.

Work package 2 has built up the offering of access to diverse high-end national Tier-1 HPC resources to European researchers within the PRACE framework. The report elucidates the work in enhancing the management of the calls, the scientific evaluations, the allocation of resources and the monitoring of progress of the DECI projects. The goal has been to build an infrastructure both sustainable and scalable as the process adapts to a larger number of participants on both the service side and the researcher side. The number of participating sites, amount of resources and the number of proposals has been growing. The work has also been motivated by the continuation of the DECI programme in PRACE-3IP and by the negotiations on a Tier-1 Optional Programme by PRACE AISBL.

This document describes in detail the progress made in the second year, discusses problems met and solved, and proposes measures to enhance the processes in the future.

During the second year, a new system for project proposal submission has been prepared, with the help of other work packages. This system should be used and refined for subsequent DECI calls. A new peer review panel has been created and used for scientific evaluations for those proposals which are not nationally reviewed.

The document also discusses expanding the range of Tier-1 offerings to community and industry support, as described in the Description of Work.

The document lists statistics on the DECI proposals received and the DECI projects accepted during the PRACE-2IP project. The success of the DECI programme among the European researchers also indicates the need of such a service by PRACE.

The work done by PRACE-2IP WP2 has enabled the provision of shared Tier-1 resources to the European research community to be incorporated into the PRACE framework. The work will be continued by PRACE-3IP WP2, based on the practices established in this project, and probably later by a PRACE Tier-1 Optional Programme.

## 1 Introduction

This deliverable presents the work done in the PRACE-2IP project work package WP2 during the second year of this project which is the final year for WP2.

The main emphasis of WP2 is to offer European researchers access to high-level Tier-1 resources provided by the PRACE partners. This goal has been fulfilled through several Tier-1 calls, or DECI (Distributed European Computing Initiative) calls, and by carefully keeping track of the needs of the ongoing DECI projects and their progress. The work has also included development of the DECI process and looking for new models to improve or expand the service. In the Description of Work (DoW), the work was initially divided into three tasks: Task 2.1 Framework for resource exchange, Task 2.2 The new DECI process, and Task 2.3 Integration of programme access and virtual communities.

Task 2.1 has concentrated on the development of the resource exchange of the DECI process. The DECI process has been evolved and has been proven to work well under the PRACE projects. To make the process even more sustainable has been one of the goals of Task 2.1.

There is an ongoing plan by PRACE AISBL to continue to provide Tier-1 collaboration and resource exchange under the framework of an Optional Programme based on the PRACE regulations, after the PRACE projects finish. The planning of a Tier-1 Optional Programme (OP) has continued and somewhat affected the work plan during the second year. The work done in Task 2.1 could be exploited in the future OP work and this is one of the ways in which sustainability of access to Tier-1 resources is foreseen.

Task 2.2 has been the major effort in WP2, launching the DECI calls for researchers, managing the evaluations, acceptance of projects, assignments to computers, and finally providing the access to researchers. The calls have been very successful, achieving high levels of response. The first scientific results of the DECI Projects have been highlighted in PRACE Digest 1/2013 [2], provided by WP3.

Task 2.3 was intended to continue and develop the virtual communities' access of DEISA on the Tier-1 level in PRACE. The concept of virtual communities is similar to the programme access in PRACE-1IP for Tier-0 resources. For a two-year project, planning for multi-year access is challenging. Nevertheless, within PRACE-2IP a consensus has not been reached over community access to Tier-1 resources. This deficiency has been overcome partly by the possibility for the communities to apply for regular DECI projects. A decision was taken by the PRACE Council to delay the discussion of Tier-1 communities' access until after the successful implementation of Tier-0 communities' access.

This document first describes the work done by WP2 during the last 12 months of WP2 and the handover to PRACE-3IP. The DECI calls, the proposals and projects are then reviewed with statistics on the submitted proposals and accepted projects.

## 2 Work Undertaken in Project Months 13–24

In September 2012, the co-leadership of the Task 2.1 was changed from Jura Tarus (CSC) to Petri Nikunen (CSC) due to the former's other duties, and according to earlier plans Petri Nikunen switched to lead the task in the second year. Wim Rijks (SURFsara) was the co-leader of the task during the second year. Leader of Task 2.2 was Chris Johnson (EPCC) and the co-leader Isabelle Dupays (IDRIS). Leader of Task 2.3 was Michael Browne (ICHEC).

Monthly video conferences of WP2 were continued during the second year of the project. They were held together with WP7 Task 7.2, and with PRACE-3IP WP2 Task 2.6. Sites participating in DECI or WP2 attend it either by video or telephone connection. The minutes of these meetings are stored in PRACE BSCW.

The PRACE-2IP All Hands Meeting was held in Paris in September 2012 in conjunction with the PRACE-3IP Kick-Off Meeting. In the WP2 session, plans especially for Task 2.1 were refined for the second year.

In Deliverable D2.2 [4] it was already reported that a tool for submission of the proposals was needed to avoid much of the manual work. WP2 was involved in defining the requirements of the Project Proposal and Reporting (PPR) tool, while the implementation was done by WP10. The tool is now being used for the DECI-11 call.

It was known that HPC-Europa project would finish at the end of 2012 and its peer review panel, Scientific Users' Selection Panel or SUSP, could not any more be used for the scientific evaluation of DECI proposals that come from such countries that are either not participating in the calls (i.e. external proposals) or that do not have a national panel. One potential solution would have been to employ the PRACE Tier-0 peer review panels for Tier-1 also but it was soon apparent that this would not be possible. So, a proposal was made to the PMO that PRACE projects and Tier-1 would establish its own panel for this task. Such a

panel was assembled in a couple of months to handle DECI-10 proposals in January 2013, partly on the basis of the former SUSP members but also having new members. PRACE partners had a possibility to propose members for the panel to ensure wide representation of countries and science areas. The panel was finally called *Non-contributing Partner Committee for Tier-1 Peer Review*. The forming of this panel was decided by PRACE MB and carried out by PMO, with support from WP2.

A joint face-to-face meeting was held on 16–17 April, 2013 in Warsaw together with WP7. Presentations and discussions in this meeting were related with the concurrent DECI call and the handover to PRACE-3IP.

The DECI-7 projects ended in October 2012, see Figure 1. The DECI-8 projects, already started in May 2012, continued to run until end of April 2013. The DECI-9 projects were started in November, 2012, and the DECI-10 projects started on May 2013, both to run for a year. Decisions on the timing of each DECI call were made by the PRACE MB only a month before each call.

The DECI-11 call and its associated work were approved during the Warsaw meeting to be under PRACE-3IP. Also it was agreed that any work after July 2013 for DECI-8, DECI-9, and DECI-10 projects would be reported to PRACE-3IP. This handover of activity is transparent to the researchers, who see a continuity of service.

PRACE 1-2-3IP all-hands meeting was held in Varna in June 2013, with also a WP2 session to report what has been done and to update plans for the future.

All the milestones of WP2 were met in time; three DECI calls (DECI-8, DECI-9, and DECI-10) were launched under the PRACE-2IP project. In addition, the DECI-7 projects, launched by PRACE-1IP, were taken care of.

## 2.1 Framework for Resource Exchange

The objective of Task 2.1 was to develop existing Tier-1 processes, and investigate alternative models. The task is split into six subtasks, for further information see Deliverable D2.2 [4]:

- T2.1.1 Develop an administrative framework for resource exchange between Tier-1 sites to support DECI calls
- T2.1.2 Analyze options for a possible exchange between Tier-0 and Tier-1 sites
- T2.1.3 Refine the pilot practices developed by PRACE-1IP towards production level practices
- T2.1.4 Create contracts, processes and exchange policies
- T2.1.5 Study new usage models for Tier-0 and Tier-1 resources
- T2.1.6 Investigate flexible resource exchange policies for load balancing between sites and over time

The main outputs from the subtasks are as follows. First, a peer review panel (“Non-contributing Partner Committee for Tier-1 Peer Review”) was formed and used for the scientific evaluation of DECI-10 proposals. Second, a proposal to have a Tier-1 preparatory access was made in order to allow scientists to access systems for porting/scaling out their applications prior a research project. However, it was not approved by the PMO, due to the lack of staffing resources for Tier-1 applications support. Third, a whitepaper about implementation of administrative tools to support DECI was prepared [6].

In addition, Task 2.1 supported WP10 in the implementation of the PPR tool, which was released for the DECI-11 call. The next six subsections discuss these and other outputs from the subtasks.

### 2.1.1 Develop an administrative framework for resource exchange between Tier-1 sites to support DECI calls (T2.1.1)

PRACE-2IP supported 4 DECI calls (DECI-7 pilot, DECI-8, DECI-9 and DECI-10). Such support continues in PRACE-3IP with the initiation of DECI-11. The above DECI calls schedule requires the support of 2 calls per year (one every 6 months). In the future it is envisaged that DECI calls will continue to operate under the umbrella of a pan European Resource Exchange programme operated in the framework of the PRACE association most probably as an optional programme or in another form. The process of initiating and executing such DECI calls requires an administrative framework for resource exchange capable of providing valid information to the stakeholders of the call management as well as the potential users. The framework should be also capable of coping with the frequency of the calls, the number of participating systems and the resources associated with those systems and the number of awarded projects.

The main functional requirements for such an administrative framework can be summarised in the following:

- Initial (before the call is issued) declaration of countries and centres (sites) participating in each call.
- Initial (before the call is issued) declaration of resources committed from sites for each call, associated with information on the technical characteristics of the resources such as system architecture, processor type, total number of cores, standard core hours provided
- Detailed declaration of resources committed for each call, including technical specification of sites just before the machine allocation to accepted projects so that efficient matching of projects to machines is performed.
- Accounting of consumed core hours by each project on each site.
- Reporting of committed and used core hours for each site on a periodic or on a per DECI call basis, so that an analysis of the *juste retour* principle is performed when necessary.

The current implantation of this administrative framework is tightly coupled with the PRACE Implementation Phase projects. Declaration of resources is being done by partners within the DoW of the projects, before each call by submitting information either to the project management office or to the DECI call manager. Declarations are being performed either using Excel or by using the PRACE wiki service.

Later in the process the systems are included in the DECI Project Management Database (DPMDB) together with information about the projects. Accounting is also performed using DBMDB while reporting is a manual process that requires gathering information from various locations.

The current process has the following shortcomings:

- There is not a clear definition of the necessary milestones and their timing, for the implementation of the administrative framework.
- There is not a single location for that the necessary information is available.
- Reporting is a manual process.

In order to setup a stable administrative framework for resource exchange, questions such as who is doing what, when something has to be done, in which sequence, what information needs to be provided need to be answered. A proposal for the technical implementation of such a resource exchange framework, via either a light weight (wiki-based) procedure, or with the creation of a new web-based tool (referred to as the System Information Database or

SIDB) or via radical changes of the DPMDDB, has been specified by the working group that was setup in the context of T2.1.1. The framework can be presented and summarised with a list of milestones that is presented below:

- **M1:** Declaration of partners' willingness to participate in the DECI programme so that their contribution is evaluated whenever necessary, i.e. per DECI call.
- **M2:** An initial declaration of resource availability for the next DECI call needs to be in place so that the feasibility of this call is evaluated
  - When: Before the decision for the initiation of the call
  - How: Submit limited site information using a web based tool
  - What: Country, Partner, Site Name, Machine Name, LINPACK [Tflops], % committed
- **M3:** Next DECI Call date is being decided – usually this is done 6 months before the public announcement.
- **M4:** Site representatives insert detailed information about the systems participating in the call.
  - When: One month before the final call announcement
  - How: Submit full site information using a web based tool
  - What: Country, Partner, Site Name, System Name, System Platform, CPU, CPU Factor, Interconnect, LINPACK [Tflops], no of cores, Total Memory, Available Storage, date available from, comment, DECI call ID, % committed, committed core-hours, comment
- **M5:** The above information is being validated and based on this the call is being announced.
- **M6:** Proposal evaluation process starts (SIDB not required)
- **M7:** Final information about available systems is being updated in web based tool.
  - When: At the start of the Technical Evaluation period
  - How: Submit final full site information using SIDB tool
  - What: Country, Partner, Site Name, System Name, System Platform, CPU, CPU Factor, Interconnect, LINPACK [Tflops], no of cores, Total Memory, Available Storage, date available from, comment, DECI call ID, % committed, committed core-hours, comment
- **M8:** Final information about available systems needs to be in DPMDDB
  - How: Import from web based tool
- **M9:** Projects are being assigned to machines
  - When: After peer review takes place
  - How: Using information available in DPMDDB
  - What: Project and machine information
- **M10:** Projects start.
- **M11:** Accounting data evaluation
  - When: Periodically
  - How: import accounting data from DPMDDB to web based tool
  - What: Core hours consumed during the reporting period per site.
- **M12:** Produce reports/views
  - When: upon request
  - How: Using the a web based tool
  - What:
    - Comparison of committed vs. used core hours for each system or partner over time periods (i.e. 1, 2, 3 years) or DECI calls.
    - Comparison of initial commitments vs. actual commitments for each system or partner over time periods (i.e. 1, 2, 3 years) or DECI calls.

In order for the framework presented above to be implemented, we propose that support from a tool (preferably web based) is needed. Due to the fact that radical changes to the DPMDDB need significant effort and major security adjustments, the development - implementation of a new web-based tool (SIDB) that would be used for the registration of sites and systems that will participate to each DECI might be needed. This tool will consolidate the actions taken by the support of different existing mechanisms i.e. email, excel, wiki. SIDB can either be implemented as a completely separate tool or a component of the DPMDDB tool. The data from this tool will be imported to DPMDDB without changing its current data structure. This would reduce the DPMDDB admin effort to import site detailed configuration data and the errors that would happen when entering / updating the respective DPMDDB information.

As an intermediate solution the use of the PRACE wiki is proposed. The use of the wiki as a tool for collecting information on Tier-1 systems has already started by the Operations activity of the PRACE-2IP project and at the moment it acts as an information point for registering and viewing information on the available systems and their characteristics for the DECI calls.

### *2.1.2 Analyze options for a possible exchange between Tier-0 and Tier-1 sites (T.2.1.2)*

Tier-0 calls and Tier-1 DECI calls are targeting different kind of proposals. To be accepted as a Tier-0 project, the investigator has to demonstrate that his application code shows a certain predefined minimum parallel scaling efficiency. Due to the size of the Tier-0 machines, very large proposals, in terms of resources can be honoured. Tier-1 machines are in general smaller, and because of the way resources on these machines are committed to PRACE the proposals that can be accepted on these machines are much smaller in size. Scalability demands for Tier-1 projects are less severe than for Tier-0 projects.

One can expect that some proposals will end up in the wrong program, for example because the scaling of their codes does not meet the minimum requirements of Tier-0 or their demand for resources is much too big for Tier-1.

The purpose of this subtask was to investigate possibilities to redirect proposals that were submitted to a Tier-0 call, but were deemed unsuited for this call, to a Tier-1 DECI call. Or, vice-versa, to redirect a proposal that was submitted to a Tier-1 DECI call, but was deemed more suited for Tier-0.

It has been argued that moving proposals from Tier-0 to Tier-1 or vice-versa is an undesired feature because the awarding criteria for the two programs are quite different and scientists know exactly what they are doing and deliberately choose one or the other programme to submit their proposal to, in order to maximize their chance of acceptance of the proposal.

This presupposes that scientists have a quite thorough knowledge of the configuration and size of Tier-0 and Tier-1 systems and how their software codes map to these systems. Without doubt this is true for a lot of scientists, but nowadays more and more scientists see HPC systems and software merely as tools to use in their research, without knowing all the ins and outs of such systems.

Furthermore, redirection only is an option if it is obvious that the proposal does not meet the criteria of the program to which it is submitted (as judged by the Technical Review). Instead of simply rejecting the proposal PRACE could offer as a service the suggestion to redirect the proposal to the other program (in time for its technical and scientific review) and make such a transfer as easy as possible.

That the investigation of possibilities to redirect proposals between the Tier-0 and Tier-1 is not just a mere theoretical exercise, is proven for example by two occurrences in the recent DECI-10 call where two proposals (MOP (Dutch) and PAMOP (Swiss)) were advised to resubmit their proposals to Tier-0, mainly because the amount of requested core-hours for these projects was an order of magnitude above the average DECI project, while the scalability was such that they were deemed possible candidates for Tier-0 projects.

What this subtask had to do was minimize the overhead and the delay of such a transfer of proposals to another call.

The situation in the first year of PRACE-2IP was such that a transfer from Tier-1 to Tier-0 (or vice versa) implied completely writing the proposal anew and waiting for half a year to resubmit the proposal, because the calls for Tier-0 access and Tier-1 access were concurrent. Therefore a proposal could not be redirected to the other Tier within the same call cycle.

### **Work in the second year**

In Deliverable D2.2 [4] an analysis of this situation was described and some recommendations were made on how to proceed with the implementation of a possible exchange process between Tier-0 and Tier-1 calls. The main conclusions and recommendations were:

1. If the possibility to transfer the proposals within the same call cycle was created, actually the synchronization of the calls for Tier-0 and Tier-1 proposals was a first step to an easy transfer of proposals between the calls, as it then helps to minimize time loss for PIs if proposals get redirected.
2. Use of the same or a similar tool for submission of the proposals, preferably with similar fields to fill in and an option to easily adapt these fields would also improve the transferability of proposals. Apart from the improvement of the management of the DECI calls this makes a redirection of proposals a lot easier because it lessens the burden of rewriting the complete proposal.
3. For a synchronized call the best point in time to redirect proposals is after the Technical Evaluation of the proposal, but before the peer review process.

In the second year of PRACE-2IP we further developed these ideas and recommendations.

For conclusion 1, it was unfortunate that from DECI-10 and the 6th Tier-0 call the period that the calls were open started to diverge. This is not a showstopper for redirection of proposals, but it will have a negative impact on the time between submission and the time of acceptance for projects that are eligible for redirection. So from the point of view of offering a relatively easy service to redirect proposals from Tier-0 to Tier-1 and vice versa it would be better to make the two calls concurrent again. If that is not possible, positioning the open period for Tier-1 calls three months after the Tier-0 call is opened, in the middle between Tier-0 calls (both have a 6 month cycle) would be the next best option as it would make the delay for misdirected proposals just as long for Tier-0 as for Tier-1 proposals.

In the second year of PRACE-2IP the intensive discussions between WP10 and the special WP2 taskforce for the introduction of a web-based submission tool resulted in a first version of such a tool for the Tier-1 DECI-11 call. This tool was based on the PPR-tool that is used in Tier-0 calls, but the design of the forms and the requested information therein is still based on the previous paper forms that were used in the older DECI calls. In future calls we will look into the possibilities for the fields to converge to the corresponding ones for Tier-0. This means that in future proposals (at least the documents) could be moved from one submission procedure to the other, without having to put much effort into it from the point of view of the owner of the proposal.

An additional benefit of using the same tool to submit the proposals, even if the fields are not the same is that one can define a single common entry point to the tool where one can publish an introductory text to describe the criteria for Tier-0 and for Tier-1 admission as clearly as possible to avoid misunderstandings about the scope and possibilities of the two different tracks and to minimize the number of proposals that are directed to the wrong call.

Even if the calls are not concurrent, the best point in time to suggest redirection of proposals is still after the Technical Evaluation, because that is the point where it should be clear that a proposal is inappropriate for the call in question. Because the scientific peer review processes of the Tier-0 and Tier-1 DECI calls differ so much, we do not foresee that it will be possible to pass on peer review reports from one call to the other, so waiting with a switch until after the peer review process would be a waste of effort from the scientific reviewers.

### **Conclusions and further work**

In conclusion we can say that by the introduction of the web based submission tool progress has been made to facilitate transferring proposals between the Tier-0 and Tier-1 DECI call program. In future DECI calls the submission form can possibly converge more to the same form as for Tier-0 so the exchange of proposals will be made easy even more.

#### *2.1.3 Refine the pilot practices developed by PRACE-1IP towards production level practices (T2.1.3)*

As reported in D2.2, DECI-7 was considered as a “pilot” DECI call intended to smooth the transition between the running of DECI under DEISA and running it under PRACE. As such, most of the work of this subtask was completed and reported on during the last reporting period. Many lessons were learnt from the pilot call and practices identified there have been employed for subsequent DECI calls. Since that time, processes have been constantly monitored and improved and DECI has increased in scope and popularity as the number of sites involved has increased and the number of applications received has increased greatly. For DECI-11 (running under PRACE-3IP) we have implemented a new submission tool for proposals and for future calls we are looking at ways of improving the DPMDB, the tool used for day-to-day administration of DECI.

#### *2.1.4 Create contracts, processes and exchange policies (T2.1.4)*

The main outcome of subtask 2.1.4 is the development of a “resource exchange agreements”. In the DoW this was planned to be ready in month 4 of PRACE-2IP as described in milestone MS21 “Agreements for DECI resources ready”. According to the DoW it should have been incorporated in deliverable D2.1.

As explained in the annual report last year, due to the decision to plan and open the DECI-7 call before the official start of the PRACE-2IP project (as a pilot call), the commitment of DECI-7 resources by partners was effected just after the start of PRACE-2IP and researchers were given access to PRACE Tier-1 resources from November 2011. Therefore the planned milestone no longer fitted with the call schedule which had been anticipated.

A first draft of the “Resource Provider Agreement” was produced in project month 6 (March 2012), and after internal review and adaptations was passed on to PMO. WP6 was involved in the discussions since they have primary responsibility for incorporating new partners and systems into the Tier-1 PRACE infrastructure. Resource Provider Agreements were introduced for the DECI-9 call. Lately, in the context of the Optional Programme, there have been further discussions about the scope and content of future contracts, processes and



exchanges policies for Tier-1. These discussions draw heavily on the work undertaken in this subtask.

### 2.1.5 *Study new usage models for Tier-0 and Tier-1 resources (T2.1.5)*

Deliverable D4.3.2 of PRACE-1IP [5] lists three access types that might be useful for the European research community: industry access, community access and preparatory access. All of them would open access to HPC resources for new users. However, the goal of PRACE is not to have as many users as possible, but to serve the European research community as a whole by enabling high-end computational research. Therefore PRACE must maintain focus, and not make too many access types available. On the basis of this, industry access to Tier-1 resources was opened as the initial priority, whereas community access and preparatory access were not.

In the next few paragraphs, we explain in more detail why only industry access was taken forward and others not. We also make recommendations on further development of industry access and DECI.

#### **Industry access**

Within the range of the second year of PRACE-2IP project, the subtask covered a good number of different options for opening PRACE Tier-1 resources to industry ranging from open R&D and pre-competitive access to preparatory Tier-1 access for industrial users. Following the decision of PRACE aisbl on early 2012 to support industrial research projects on Tiers-0 systems through Open R&D (with publication of the results at the end of the grant period) the subtask group gathered around the recommendation that at this stage of PRACE, an open access for industry is a viable solution because there would not be sense to prevent industrial research projects from using PRACE Tier-1 resources if the projects are as good as academic projects.

In the DECI-9 call, the scope of the call was then extended to industrial projects. Such research projects must be Europe-based and contain non-proprietary research that will be published in peer-reviewed journals. In that regard, the only thing that differentiates an academic proposal from an industrial proposal is the principal investigator's place of work. Also some academic proposals contain industry-related research.

Two industrial pilot projects applied for resources in the DECI-9 call. Both were accepted (ESM4OED and FORSQUALL). The industrial proposals went through the same application process as other proposals, starting from proposal submission, and ending to scientific and technical evaluations. Because some countries have restrictions against non-academic research, industrial projects are supported from the external projects' budget – in the same way as academic projects that come from countries not participating to PRACE. DECI-10 and DECI-11 calls did not draw any proposals from industry.

Industry access on Tier-1 systems will be continued, either within DECI and/or through a new separate evangelisation call prepared by PRACE-3IP WP5. In future, it is worth investigating an option of special hybrid joint proposals that are led by PI from the academia but the proposals must also include industrial partners including large companies as well as SME collaboration or identify potential SME candidates that could be evangelised for using HPC in order to increase their competitiveness.

In DECI-9, scientific review of the industry projects was done by national evaluation committees. This may be problematic because the resources are allocated from the external projects budget. National panels may overrate industrial proposals coming from the same

country because of national interests. Therefore we would recommend using the “Non-Contributing Partner Committee for Tier-1 Peer Review” for this purpose.

### **Community access**

Scientific communities can apply for PRACE Tier-1 resources through DECI, but there are no separate community calls for them. The question whether such an opportunity should be opened was discussed in the subtask. The decision was not to recommend it. There are two reasons. First, the number of access types should be kept small. Second, a separate community access might in the worst case put DECI projects and community projects in unequal positions in resource allocation. Better would be to notice the needs of scientific communities within DECI.

Scientific communities have three properties that are less common, but existent, among DECI projects. First, scientific communities are multinational. To increase the probability of a multinational project to get resources, those projects whose country of origin is questionable could be supported from the external projects’ budget, and increasing the external projects budget accordingly. An alternative solution would be to support multinational projects from joint country budgets based on collaborators’ place of work.

Second, scientific communities last longer than a typical DECI project, and they need resources over this longer period. This requirement is difficult to take into account in the DECI process without treating projects unequally. However, an expectation for a scientific community to get resources is high, so the only thing that the scientific community needs to do currently to get resources is to send a new DECI proposal once a year. This should not be a problem for them.

Third, scientific communities would appreciate a possibility to sub-allocate resources given to them. This wish is not met in the Tier-0 Multi-year Access (or other Tier-0 access types), because it could potentially compromise peer review principles of PRACE. For the same reason, this opportunity is not recommended for DECI either. Besides, allowing resource redistribution to all DECI projects would cause a lot of extra administrative work.

### **Preparatory access**

There are user groups, especially in industry that would benefit from preparatory access to Tier-1 systems. Preparatory access would harmonise Tier-0 and Tier-1 processes and make DECI projects readier for computation. That would help to even out the load that DECI projects currently put to computer systems: the usage is low in the first months of execution and increases gradually towards the end.

A proposal to have a Tier-1 preparatory access call with the same time as the DECI-11 call was made and passed to the PMO, but it was rejected on the basis that it could be mixed up with the Tier-0 preparatory access and that resources for Tier-1 applications enabling were limited. In the future having in mind that PRACE sets up SHAPE (SME HPC Adoption Programme in Europe), that aims at raising awareness and provide SME’s with the expertise necessary to take advantage of the innovation created by HPC, the preparatory access may be reconsidered once again. For the preparatory access proposal the requirements of a form were defined, and they are included in this Deliverable as Annex 9.1.

#### *2.1.6 Investigate flexible resource exchange policies for load balancing between sites and over time (T2.1.6)*

The main objective of this subtask was to investigate whether there is a need to improve the load balancing of resources over sites within the same DECI call and if there is a need for

resource exchange policies that span multiple DECI calls. If such a need is there, a second objective of this subtask was to suggest ways to implement this.

From the start of the DECI project resource exchange was based on the *juste retour* principle, the idea that each site could claim as much project resources as it committed to the DECI call (minus a small percentage that was donated to external projects of sufficient scientific quality). The *juste retour* principle was applied per call. So for each new call a balance was sought between committed resources and consumed resources per site and no thought was given to the idea to extend the concept over multiple calls. There was little reason to think about this, because each call was heavily oversubscribed by all partners, so there was no room to build up debts or credits. If a site wanted to allow extra project resources, they had to compensate these immediately by committing extra resources themselves, thereby keeping the balance intact.

In the first year the results of the discussions in PRACE-1IP WP4.3 [5] were analyzed and some bottlenecks in this task were identified.

As stated in Deliverable D2.2 the main topics that should be tackled in the second year were on the process of expressing all resources in standard core-hours and on the procedures to match commitments and project claims per partner per call. To that end two small taskforces were created, one to discuss conversion factors and one to discuss the match between commitments and claims.

The results of these discussions were some conclusions and some recommendations that are documented here. Some of the recommendations are already accepted by WP2 staff and implemented. Implementation of some other recommendations will be continued in PRACE-3IP.

### **Matching commitments and project resource claims per call**

In the first instance, matching commitments and project resource claims is a responsibility of each partner itself, because it is the local DAAC (DECI Access and Allocations Committee) that ultimately recommends which local projects to accept and to ensure that the resources for these projects are matched with the local commitments.

The second step is to confirm the approval of the local projects in a general videoconference by DAAC. In this videoconference the acceptance of external projects is also decided, based on the final ranking done by the newly appointed “Non-contributing Partner Committee for Tier-1 Peer Review” and DECI staff. The amount of resources that is available for these external projects has to be agreed upon before the start of each DECI call. Traditionally it is 15% of the nominally committed resources.

If after this process there are still resources left over, in a follow up step partners can submit claims for extra resources for their local projects (so without matching these locally) and the assembly can decide to award extra resources to external projects if these have sufficient potential for excellence.

The concept that there are excess resources after the second step is fairly new to DECI and is due to the several facts:

- The number of partners is expanding from the original 11 DEISA partners and the new partners often do commit resources to a DECI call, but do not have always sufficient good proposals to match these.

Existing partners are expanding or renewing their Tier-1 machines and thereby contribute more to DECI.

- The frequency of DECI calls has increased to twice per year. As a consequence for some partners the resource claim of proposals per call is sometimes below their nominal commitment.

In a next step the accepted proposals have to be allocated to the participating sites, in such a way that all committed resources can indeed be consumed by the projects. This is a collaboration of T2.2 and T7.2 and involves the matching of the preferences for certain machines as expressed in the proposals to the available resources from the commitment of all the partners. Of course this is never a perfect fit, so partners that are slightly overcommitted have to confirm that they accept this over commitment.

Currently, both the allocation of leftover resources to additional projects and the over commitment of resources at certain sites is done voluntarily, without acquiring any debts or credits.

A trend we see in recent DECI calls is that the number of external projects is growing rapidly. This is of course a good development as it demonstrates that the Tier-1 DECI program has embedded itself in the European research community. But it raises the question if we have to increase the amount of resources that are made available for these external projects (traditionally it is 15% of the nominal commitments of all partners). Our recommendation would be to not increase the share for external projects formally, but to reserve some excess resources (if these are present) to accept additional external projects.

### **Matching commitments and project resource claims over multiple calls**

Now that we sometimes encounter an excess of resources in the more recent DECI calls, it is time to discuss whether we have to handle this phenomenon formally or if we proceed as we have done so far. Up until now we treat each DECI call as if all partners participate with a clean slate without carrying any history from previous calls. All excess claims or under allocations are regarded per call and are handled on basis of mutual agreement without any formal contracts of any kind. This process seems to be working well as there are no discussions yet about unfair distribution of resources. Several sites even allow for over allocation at the end of projects, if the PI provides good arguments for this (disappointing performance of applications, running some extra simulations to gain extra insight), demonstrating that most sites are not overly concerned with an exact match of commitments and claims.

Our suggestion would be to leave this procedure as is as long as no disputes arise over this question. The alternative would be to set up an administration of debts and credits for each site and base future allocations and grants on the net sum of these two factors. A disadvantage of this approach is that sites that are continuously delivering a lot of good proposals that are honoured in the DAAC sessions, build up debts quickly that would be difficult to compensate for in future calls. This could possibly be remedied by taking other resources (long term storage for example) into account, but this would further complicate things because you have to come up with an accounting unit that accounts both compute resources as well as the additional resources. That would mean introducing a lot of bureaucracy to remedy something that has not been a problem up until now.

### **Measuring tools and Conversion factors**

To be able to check if consumed resources are matched with commitments certain bookkeeping procedures are necessary. To this end we use DART (Distributed Accounting Report Tool), where most sites automatically register the monthly usage numbers for DECI projects. These numbers are then extracted and imported into the DPMDDB, which is the central tool to manage the DECI projects.

To compare the consumed resources on different machines one needs a translation from the local accounting unit to a common DECI accounting unit. To this end the “conversion factor” was introduced in the early days of the DECI program. From the start of DECI in the DEISA project the concept of “conversion factor” was used for two things:

1. To express all the requested resources from all proposals in a common unit
2. To express the relative performance of all architectures which participate in a DECI call, in that same common unit.

For historical reasons the unit is chosen to be 1 core-hour on an IBM Power4+ (1.7 GHz) processor.

Ideally the conversion factors should be determined by running extensive benchmarks on all different architectures and compare the results to the same benchmark results executed on the Power4+ processor. The conversion factor would then be a weighted average of the speedup factors of all the different benchmarks. This would be a very time consuming business and after phasing out the original reference processor it would become difficult to adapt the benchmark suite to new developments in software improvement and change of favourite applications for certain type of simulations, because these newly introduced codes could not be run on the reference processor anymore. Furthermore an averaged performance figure is only a coarse estimate for the relative performance of specific applications. So there is no point in being overly accurate in the determination of the conversion factor.

So it was deemed that the relative importance of obtaining correct conversion factors did not warrant all this effort. Therefore from the introduction of the conversion factor a more pragmatic solution was adapted. The conversion factors were determined by a small (unofficial) panel from DECI staff and the site hosting the machine in question on a best effort basis. The chosen conversion factor was deemed to be a good estimate if the hosting site was satisfied about the “return of investment” for the committed resources and if the scientists were happy about the amount of work they could get out of their allocations. So far this has worked to the satisfaction of all partners involved. Therefore we propose not to change this concept but only to improve on the procedure.

One improvement could be to make the panel on conversion factors an official part of the procedure. This proposal was accepted on the WP2 all hands meeting in Warsaw, in April 2013 and the first members of the panel were accepted at that meeting.

Another improvement is to document the rules of thumb and the considerations that led to the determination of the specific conversion factors, so they can be used as basis or reference by the determination of new conversion factors. To this end we created a page on the internal Wiki, where these considerations can be documented for reference by DECI staff. All staff members can update this page, both on conversion factors for participating machines, as well as for machines that are mentioned in proposals as basis for the resource request, but that are not part of the Tier-1 infrastructure.

Once the conversion factor is determined it is imported into the DPMDB. Here it is used to translate the monthly usage numbers, expressed in local units, into the common DECI unit. The DPMDB already contains features to produce an overview of the usage numbers per project per month and per site. These overviews are then used in the monthly videoconferences where progress of the projects is discussed and the matching of claims and commitments is monitored.

DART and the DPMDB already contain a lot of information and have a lot of features to display all kinds of views on the data contained in it. However, there are still a number of improvements that could be made to it. For example:

- All sites have not yet hooked up with DART, although the vast majority have.
- Extracting monthly usage numbers from DART and importing them in the DPMDDB is still done manually.
- Although overviews of commitments and claims spanning several DECI calls can already be extracted from the DPMDDB with some effort, features to easily produce them on the spot are not yet present in the DPMDDB. These overviews can become handy if disputes arise over the fair distribution of commitments and claims over sites or to produce statistics on the use of the Tier-1 infrastructure.
- Adding new features to the DPMDDB is difficult at the moment. Therefore it is desirable to restructure the DPMDDB to facilitate the adding of new features.

Therefore WP2 requested assistance from WP10 (to be continued by WP6.3 in PRACE-3IP) to improve upon the administrative tools that are used to keep track of the usage of Tier-1 resources. The current status is that WP10 together with WP2 are taking inventory of what features should be added and how to implement these.

## 2.2 DECI process

The past year has seen the ending of DECI-7 projects, the continuation of DECI-8 projects, the starting of DECI-9 projects and the opening of the DECI-10 call followed by the starting of projects and the opening of the DECI-11 call. DECI-11 will be considered completely under PRACE-3IP. Figure 1 shows the timeline of DECI calls. In the following sections we give an overview of the DECI calls involved in this reporting period. More detail about the statistics of proposals and accepted projects can be found in Sections 2.2.1–2.2.4 and 5–7 below.




Figure 1 Timeline of DECI calls showing those dealt with under PRACE-2IP and PRACE-3IP

### 2.2.1 DECI-7

Projects accepted for the DECI-7 “Pilot” call were given until 31st October, 2012 to finish production runs and then a further 3 months to remove all data and produce appropriate reports. These reports will be dealt with by WP7, T7.2. Several projects were given short extensions due to some delays earlier on in the project but all have now finished. In DECI-7, 54 proposals were received and 35 projects were accepted.

### 2.2.2 DECI-8

Projects accepted for the DECI-8 call were given until 30th April, 2013 to finish production runs with the same arrangements for data removal and reporting as for DECI-7. Some projects have given extensions but almost all are finished now. Management of DECI-8 was performed via monthly video conferencing and using appropriate DECI tools such as the DPMDB and the BSCW. In DECI-8, 49 proposals were received and 33 projects were accepted.

### 2.2.3 DECI-9

Technical and Scientific evaluations of the 45 DECI-9 proposals were performed as for previous DECI-8 and 31 projects were accepted. The production run phase started on 1st November, 2013 and will last for 1 year. Management continues as for previous DECI-8.

### 2.2.4 DECI-10

The DECI-10 call opened on 5th November, 2012 and closed on 14th December, 2012. 88 proposals were received but as 3 of them arrived late 85 were taken forward for consideration. Technical evaluations were performed as for previous DECI-8.

Scientific evaluations for the internal proposals were performed as usual. However, for the external proposals it was not possible to use the HPC-Europa SUSP as the HPC-Europa project finished on 31st December, 2012. To deal with this, a new panel was organised exclusively for the purposes of reviewing DECI proposals. As had happened with the SUSP, the new panel (Non-contributing Partner Committee for Tier-1 Peer Review) was set up to review all external proposals and any internal proposals which could not be reviewed within the country of origin, for example, for sites which were new to DECI or for countries which had too few proposals for a ranking to be meaningful. Candidates for the new panel were put forward by partners via the PRACE Management Board, which made the decision (PRACE-2IP-MB-05.d3) on establishing the panel in the meeting on 18<sup>th</sup> December, 2012. 33 reviewers were selected across 13 countries and 28 of these took part in the reviewing of the 34 DECI-10 proposals. After most reviews were completed a chair was selected and a ranking meeting was convened. Each proposal requires 2 reviews meaning 68 reviews were sought in total. This number of reviews put a lot of stress on the administration of panel and not all reviews were received in time for projects to start on 1st May, 2013.

In light of this some changes will be made for DECI-11 such as the use of a new submission tool for the panel to access and complete their reviews and we also intend some improvement of the evaluation form itself. In addition we will do what we can to minimise the number of reviews given to the panel by trying to arrange for contributing partners to perform reviews for internal proposals within their own countries. In total 37 projects were accepted for DECI-



10 and the production run phase started on 1st May. Any projects which were delayed in starting due to the above problems will be given an appropriate extension wherever possible. Management continued as for previous DECI. DECI-10 will be considered under PRACE-2IP until 31st July 2013, whereupon all DECI work will transfer to PRACE-3IP.

### 2.2.5 *DECI-11*

DECI-11 was considered under PRACE-3IP right from the beginning of preparations to open the call.

### 2.2.6 *DPM Database*

The DPM database (DPMDB) continues to be the main tool used for running the DECI projects where proposal evaluation, links to proposal, evaluation forms and final reports, accounting information and project progress are all recorded. The DPMDB will continue to be used for the foreseeable future but it is becoming abundantly clear that major work is needed to improve and update the tool. Discussion with PRACE-2IP WP10 and PRACE-3IP WP6 are presently underway to facilitate these improvements.

### 2.2.7 *Scientific evaluation forms*

The Scientific evaluation forms used were not changed between DECI-7 and DECI-10 and evaluators were asked to assess the following points:

- How well the proposal matches the call
- The quality of the proposal
- The likely impact of the proposed work
- The applicant's ability in terms of his/her track record and skills
- The resources requested

However, we intend to make some improvements to the form for DECI-11. For example, it would be better to ask the reviewers how suitable the proposal is for receiving Tier-1 resources rather than asking how well the proposal matches the call and it is not easy to assess the track record and skills of applicants as this is not explicitly asked for in the proposal form.

## 2.3 **Community Access**

In the context of PRACE, Community Access has three primary aims. First it allocates resources based on defined criteria at a community level rather than a project level. The sub allocation of these resources to individual researchers is then delegated to the community itself. Second it provides for longer-term access than traditional project access which is typically one year. Additionally, it could provide for a form of blended access spanning national and PRACE access.

Further to comments in section 2.1.5, which amongst other things discusses the limits of current access for communities, this section addresses the concept of a broader and integrated provision of service. Community access in general will desire significant resources perhaps equivalent to that of several project style requests. Furthermore, continuity of access over extended periods can be crucial. For example the access may be requested in parallel with a physical instrument or another multiyear funding programme. It was decided by PRACE MB before the beginning of PRACE-2IP that a pilot community access call would not be opened. PRACE has a justifiable concern that it may find itself unable to honour promises made to

communities if there was a gap in service provision while PRACE transitions from current to future models of operation.

Given this significant restriction the original core mandate of this task as per the DoW to “provide Tier-1 Resources for virtual communities and EU projects” cannot meaningfully be fulfilled at this point and hence is addressed by means of this short report. However that is not to say that the original mandate should not be incorporated into future initiatives when the appropriate access programmes becomes a viable option. At which point the requirements should be reevaluated in light of community targeted efforts elsewhere in PRACE but also changes elsewhere in the European research ecosystem e.g. FET<sup>1</sup> projects. For example the Human Brain project<sup>2</sup> or the Graphene Flagship project<sup>3</sup>, both of which have a significant and broad need for computational access. Furthermore, should Centres of Excellence, as referred to in the Commission’s communication on HPC [7], be developed these too will require diverse computational access on a long-term basis.

Throughout the course of PRACE-2IP considerable discussion on this topic took place but this is not commensurate with the level of work initially planned to implement a community access programme. Much of the effort has been redeployed to the work package, largely to the benefit of researcher community via management of the Tier-1 DECI access programme.

This delay has not prevented existing community researchers from applying to access programmes both at Tier-0 and Tier-1, which have expanded significantly over time, making available resources to communities albeit in a project rather than community model. Access at the project level remains free of charge and based on scientific excellence. Work of the PRACE-1IP project (work package 7) has strengthened the position of communities when applying for such project access by enabling the applications. This has been complemented in PRACE-2IP by the work of work packages 7 and 8. Examples of communities that have benefitted directly through the work of other work packages to date are: IS-ENES (climate), MAPPER (multiscale modelling), Scalalife (life sciences), Vercé EPOS (earth sciences). A great deal of work has been carried out on very many independent community codes, e.g. Quantum ESPRESSO and Dalton.

Furthermore the very many white papers, including 40 specifically on application scalability, produced by PRACE form a resource for all communities with an interest in HPC.

Engagement with scientific communities will form part of the future evolution of the PRACE Infrastructure beyond the current PRACE Implementation Phase projects. These future phases of PRACE’s evolution aim to put in place sustainable structures for the Research Infrastructure as a whole. Through working with communities, we are developing a better understanding of their future requirements. This means that, when we are in a position to offer community access, and we will be able to design and deliver access schemes which fit their needs. PRACE’s cognisance of the needs of communities is further expressed in the PRACE Scientific Case [8], a documented prepared together with community representatives.

With this future activity in mind it is prudent to note that the earlier engagement between PRACE and communities is initiated at a high level (i.e. above the level of the individual researchers or PRACE work package leadership) the better, as this allows concerns to be factored into future negotiations to develop a suitable joint position and understanding.

---

<sup>1</sup> [http://cordis.europa.eu/fp7/ict/programme/fet/flagship/home\\_en](http://cordis.europa.eu/fp7/ict/programme/fet/flagship/home_en).

<sup>2</sup> <http://www.humanbrainproject.eu/>

<sup>3</sup> <http://www.graphene-flagship.eu/GF/index.php>

### 3 Collaboration with PRACE-2IP WP7 Task 7.2

In the second year of PRACE-2IP, WP2 continued close collaborations with task T7.2 in WP7 (Applications Support for new DECI Projects) during the whole DECI process.

As agreed at the DECI sessions of PRACE-2IP All Hands Meeting in Paris (September 2012), the following were carried out during the second year of PRACE-2IP, to improve the collaboration between WP2 and T7.2-2IP.

- Technical Evaluations (TEs)

TEs were completed by T7.2-2IP. The completed TEs were then passed to WP2 for the next step Scientific Evaluations. In the second year of PRACE-2IP, the TEs for DECI-9 and DECI-10 were completed for the received DECI applications.

- System Assignments

The efforts for the system assignments were moved from WP2 to T7.2-2IP. The system assignments are based on the TEs to assign the most suitable architectures to the accepted DECI projects where possible.

- Reduce Overlap between WP2 and T7.2

To reduce the overlap of DECI projects management between WP2 and T7.2, it has been agreed that WP2 will be responsible for the general DECI projects' progress tracking and T7.2 will focus more on the applications enabling support which required more than 1 month in their original proposals.

T7.2-2IP is responsible for helping the DECI users with getting access to the assigned Tier-1 systems and the initial meetings with PIs at the starting stage of each DECI phase. WP2 will be responsible for the progress tracking of all the ongoing DECI projects. If there is any issue raised, WP2 will then pass the questions to T7.2 and T7.2 will provide technical support where needed.

Besides, the DECI sessions at the PRACE All Hands Meeting in Paris, there was also a WP2/WP7 joint DECI session at the WP7 face-to-face meeting in Warsaw (April, 2013). Several topics were discussed by the WP2 and T7.2 members together, including the allocation scheduling policies across sites, the new submission system (PPR tool), the maintenance and improvement of existing DECI tools (e.g. DPMDB, DART, Workflow portal, etc). The WP2/WP7 joint DECI session had a very productive discussion to improve the DECI work further and similar joint DECI sessions will probably be planned in the future WP7 face-to-face meetings.

### 4 Collaboration with PRACE-3IP WP2

The lifecycles of DECI-9, DECI-10 and DECI-11 run through a period of overlap between PRACE-2IP and PRACE-3IP. To deal with this DECI-9 and DECI-10 will be worked on by PRACE-2IP up until the end of July 2013 and DECI-11 will be worked on completely by PRACE-3IP. This is indicated by a dotted line in Figure 1. Most of the procedures and management of DECI will move over seamlessly from PRACE-2IP to PRACE-3IP. Early preparation work for the new submission tool was done under PRACE-2IP with a view of implementing this for future DECI calls. The implementation of this was done under PRACE-3IP for DECI-11.

## 5 Overview of DECI-9 Projects

### 5.1.1 Projects by application area

Scientific discipline	Proposals received	std-hours requested	Proposals accepted	std-hours granted
Applied Mathematics	1	25,200,000	1	5,045,759
Astrophysics	5	32,652,500	4	17,802,497
Bio Sciences	12	61,891,000	10	31,426,259
Earth Sciences	2	10,085,000	1	960,000
Engineering	2	4,745,000	2	4,276,250
Materials Science	19	102,787,235	11	38,812,045
Plasma & Particle Physics	4	66,931,300	2	6,013,800
<b>Total</b>	<b>45</b>	<b>304,292,035</b>	<b>31</b>	<b>104,336,610</b>

Table 1 Proposals received for DECI-9 together with proposals accepted

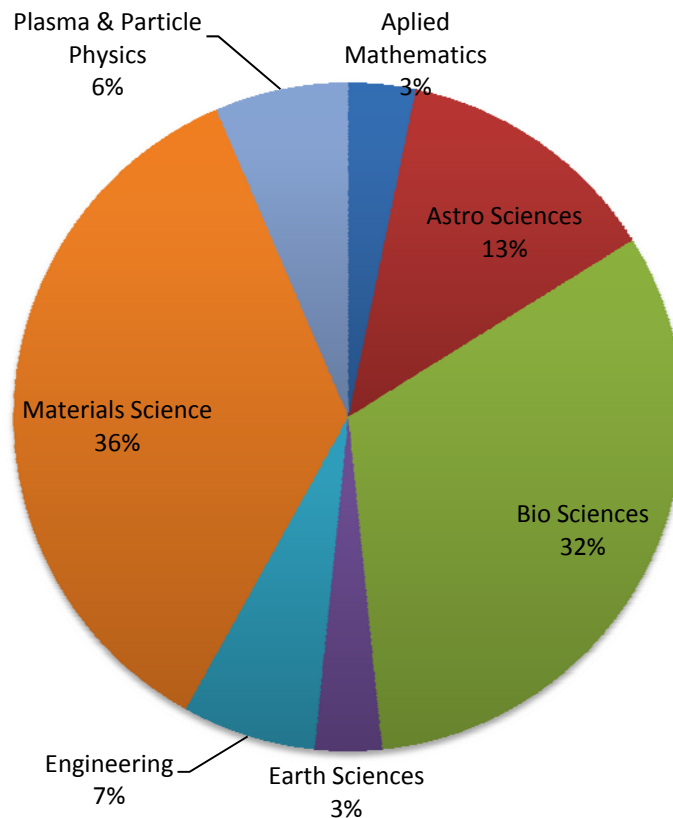


Figure 2 DECI-9 accepted projects shown by scientific discipline

### 5.1.2 Projects by country of origin

Country of PI	Proposals	Projects
Argentina	1	1
Bulgaria	1	1
Czech Republic	2	2

Denmark	1	1
Finland	2	2
France	5	4
Germany	6	4
Greece	2	1
India	0	0
Ireland	3	3
Italy	3	3
Japan	0	0
Norway	0	0
Portugal	2	1
Slovak Republic	0	0
Spain	5	2
Sweden	5	3
Switzerland	3	1
The Netherlands	0	0
Turkey	2	1
UK	2	1
USA	0	0
<b>Total</b>	<b>45</b>	<b>31</b>

Table 2 DECI-9 Proposals and Projects by country of origin

## 5.1.3 Resources Allocated

Internal/External/ Industry	DECI project	Computational resources awarded (core hours)	DECI home site	DECI execution site(s)
External	COIMBRALATT	5,265,000	BSC	PDC
External	NMRCONF	2,080,000	CSC	CSCS
External	TB-Drugs-In_silico	736,000	FZJ	CINECA
External	SPSC	7,200,000	SURFSARA	FZJ
Industry	FORSQUALL	921,600	CINES	EPCC
Industry	ESM4OED	3,000,000	EPCC	EPCC
Internal	ICREIMUTANTS	1,100,000	BSC	EPCC
Internal	SpEcBNS	2,500,000	BSC	RZG
Internal	DOPE	3,000,000	CINECA	CINECA
Internal	GPCR4D	660,000	CINECA	EPCC
Internal	iMIG	2,500,250	CINECA	BSC
Internal	AuPd-Seg	640,000	CINES	RZG
Internal	IONGATE	1,200,000	CINES	UIO
Internal	NPR-LQCD	748,800	CINES	CINECA
Internal	CompSym	5,600,000	CSC	CSCS
Internal	Planck-LFI2	7,000,000	CSC	CSC
Internal	LCRR	9,900,000	CSCS	CSCS,STFC
Internal	LBSCOM	3,200,000	EPCC	CSC
Internal	MoMoGal	7,002,500	FZJ	EPCC,UIO

Internal	Reactive_Ceria	1,500,000	ICHEC	FZJ
Internal	Si-Interfaces	2,488,320	ICHEC	CINES
Internal	SPH-WEC	1,620,000	ICHEC	CSCS,NCSA
Internal	AIMD-PAF	231,000	NCSA	EPCC
Internal	CoStAFuM	9,687,608	PDC	RZG,UIO
Internal	DifVib	6,250,000	PDC	EPCC,PDC
Internal	HydFoEn	2,500,000	PDC	UHEM
Internal	GanDaLF	1,299,999	RZG	UIO
Internal	PTACRB	4,649,997	RZG	ICHEC
Internal	HiSSor	5,000,000	UHEM	STFC,RZG
Internal	MPI-FETI	2,656,250	VSB-TUO	EPCC
Internal	NPT_MC	2,352,000	VSB-TUO	UHEM
	<b>Total</b>	<b>104,489,324</b>		

Table 3 DECI-9 Projects by home site showing execution site and resources awarded

## 6 Overview of DECI-10 Proposals and Projects

### 6.1.1 Projects by application area

Scientific discipline	Proposals received	std-hours requested	Proposals accepted	std-hours granted
Astro Sciences	9	78,194,208	6	36,301,536
Bio Sciences	13	85,509,038	7	31,823,231
Earth Sciences	8	77,855,000	6	23,390,000
Engineering	16	100,078,400	8	37,527,501
Informatics	2	2,247,153	0	0
Materials Science	36	183,525,459	10	29,284,120
Plasma & Particle Physics	1	7,500,000	0	0
<b>Total</b>	<b>85</b>	<b>534,909,258</b>	<b>37</b>	<b>158,326,388</b>

Table 4 Proposals received for DECI-10 together with proposals accepted

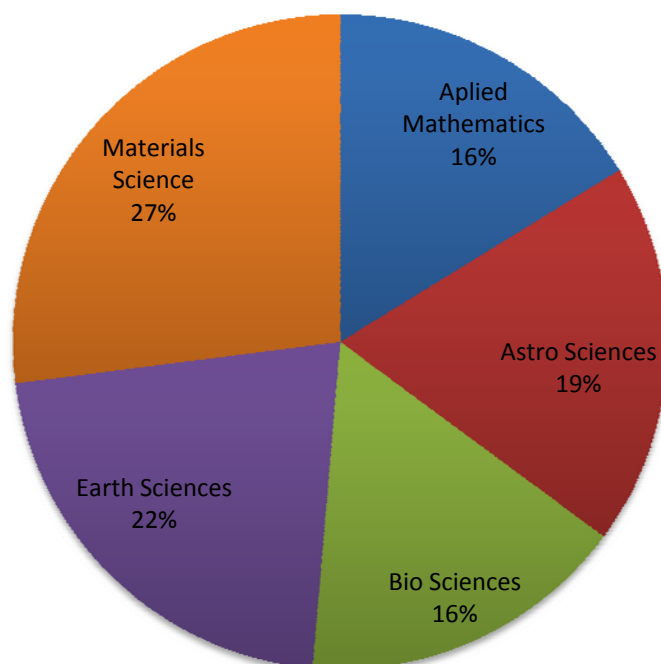


Figure 3 DECI-10 projects shown by percentage accepted by scientific discipline

#### 6.1.2 Projects by country of origin

Country of PI	Proposals	Projects
Austria	1	0
Belgium	5	1
Czech Republic	1	1
Denmark	1	1
Finland	4	3
France	1	0
Germany	8	5
Greece	7	0
Ireland	3	3
Italy	12	2
Poland	1	1
Portugal	3	2
Spain	11	2
Sweden	7	5
Switzerland	1	0
The Netherlands	3	2
Turkey	2	2
UK	14	7
<b>Total</b>	<b>85</b>	<b>37</b>

Table 5 DECI-10 Proposals and Projects by country of origin

## 6.1.3 Resources Allocated

Internal/External	DECI Project	Computational resources awarded (core hours)	DECI Home site	DECI Execution site(s)
External	AIDMP	1,200,000	BSC	FZJ
External	fplb	4,216,934	BSC	WCSS
External	NANODROPS	10,141,200	CSC	EPCC
External	Novel_Anticoagulants	480,000	BSC	CINECA
External	SPAITAC	5,625,000	BSC	EPCC
External	TheoMoMuLaM	2,835,000	SARA	UIO
Internal	APOP20X3	480,000	ICHEC	PSNC
Internal	CELESTE	5,500,000	ICM	CSCS
Internal	CONVDYN13	5,419,008	CSC	RZG
Internal	Dissipative_Phenomena	686,400	EPCC	EPCC
Internal	DIVI	3,500,000	SARA	ICHEC
Internal	DNSTF	8,437,500	PDC	EPCC
Internal	ERPP	5,910,000	CINECA	CYFRONET
Internal	EXC-XMCD	2,000,000	VSB	PSNC
Internal	GalChem	5,940,000	EPCC	SURFSARA
Internal	Galsim	10,692,000	EPCC	CSC
Internal	GREENLIGNITE	2,240,000	UHEM	VSB-TUO
Internal	HIGHERFLY	5,000,000	EPCC	EPCC
Internal	HIV1-GSL	7,000,000	CSC	EPCC
Internal	HYDRAD	6,002,196	RZG	VSB-TUO
Internal	HyVaMPI	3,500,000	CSC	UIO
Internal	INPHARMA	750,000	FZJ	EPCC
Internal	InterDef	3,739,250	EPCC	SURFSARA
Internal	JOSEFINA	6,250,000	EPCC	PDC
Internal	LargeRB2013	6,817,501	FZJ	EPCC
Internal	LipoSim	8,750,000	PDC	PDC
Internal	MEGAREACT	1,500,000	PDC	UIO
Internal	MoDSS	577,500	FZJ	RZG
Internal	MOTUS	2,625,000	CINECA	ICM
Internal	PLANETESIM-2	7,500,000	PDC	FZJ
Internal	PTACRB-2	2,950,000	RZG	CYFRONET
Internal	RODCS	3,750,000	ICHEC	CSCS
Internal	SCosPtS	2,280,000	SARA	CSC
Internal	TransMem	5,200,000	VSB	SURFSARA
Internal	waveclim	2,295,000	ICHEC	CSCS
Internal	WIND-FORECAST	280,000	UHEM	CYFRONET
Internal	WISER	6,000,000	EPCC	EPCC
	<b>Total</b>	<b>158,069,489</b>		

Table 6 DECI-10 Projects by home site showing execution site and resources awarded



## 7 Overview of DECI-11

For DECI-11, 117 proposals were received from 24 different countries of which 2 were rejected as ineligible and one has since withdrawn leaving 114 proposals for consideration. DECI-11 will be worked on entirely under PRACE-3IP and more details will be reported in the relevant deliverables there.

### 7.1.1 Overview of DECI proposals for all calls

Figure 4 shows the number of proposals received over all DECI calls. Although the pattern and number of proposals received for the calls under PRACE seems similar to that of DEISA, it is worth noting that under PRACE calls happen twice yearly whereas under DEISA they happened yearly. Taking this into account the number of proposals received for DECI-10 and DECI-11 is a very large increase on what we have previously received.




Figure 4 Number of proposals received for all DECI calls so far under both DEISA and PRACE

## 8 Analysis of Tier-1 Resources

Country	DECI-7		DECI-8		DECI-9		DECI-10	
	Granted (k std hrs)	Contributed	Granted	Contributed	Granted	Contributed	Granted	Contributed
Bulgaria	900	900			400	300		
Finland	7,400	7,600	3,200	3,200	10,200	13,600	13,000	14,400
France	7,200	7,300	5,000	4,900	2,500	2,400		
Germany	18,800	19,000	20,400	20,400	14,500	14,700	14,700	14,700
Ireland	4,500	4,500	4,300	4,500	4,700	4,500	3,500	4,500
Italy	6,400	6,400	6,100	6,400	4,200	6,400	500	6,400
Norway					10,300	12,200	7,800	12,200
Poland	9,500	9,500	4,100	4,100		5,800	18,500	26,400
Spain	1,500	1,500			1,800	2,400	-	-
Sweden	16,300	15,900	16,000	15,900	13,300	15,900	15,000	15,900
Switzerland			11,600	11,300	8,900	11,300	11,500	11,300
The Netherlands	8,200	8,200	2,500	2,600		2,600	14,900	15,700
The Czech Republic							8,200	8,000
Turkey			5,100	6,700	4,900	5,000		3,200
UK	9,800	9,800	20,400	19,700	27,900	41,000	50,400	41,000
<b>Total</b>	<b>90,600</b>	<b>90,800</b>	<b>98,700</b>	<b>99,800</b>	<b>103,500</b>	<b>138,200</b>	<b>158,000</b>	<b>173,800</b>

Table 7 Resources contributed by partner in thousands of standard hours

Table 7 shows the resources contributed by each country for each DECI. As can be seen from the table and Figure 5 the amount of resources available has risen steadily with each call and the number of countries contributing has remained fairly constant with 11 countries contributing for DECI-7 and DECI-8, 14 contributing for DECI-9 and 12 contributing for DECI-10. This is due to the fact that some sites do not contribute for every call (e.g. BSC, Spain contribute to alternate calls) some new sites join each year and others have dropped out, such as CINES and IDRIS in France.




Figure 5 Resources contributed per DECI call

## 9 Annex

### 9.1 Tier-1 preparatory access form requirements

In the context of the work being done within WP2 for the management of the DECI calls and the planning for possible extension of such calls in the future, the idea of supporting preparatory access to the Tier-1 systems was evaluated.

The evaluation included the identification of the need and modalities of such an access methodology as well as the definition of additional requirements mainly in the area of the design of the application form for such projects.

Bellow we present some of the main findings:

- From the point of view of the DECI management teams, preparatory access to Tier-1 systems could be a useful service for the European scientific and industrial communities. Examples of such usage are given bellow:
  - By scientists and researchers that are using codes not yet tested on new hardware architectures, i.e., many core systems, systems with accelerators etc.
  - It could also potentially be of use to the scientists and researchers that are developing new promising codes and have limited resources available to them to test scalability.
- The available computational resources for such projects would be relatively low and adequate for simple development and scalability testing.
- Emphasis should be given to porting, optimization and scalability testing of such applications.

To enable users propose such applications it was deemed necessary that some changes in the proposal submission form to be done so that the specific requirements can be captured and the form is simplified to attract more users.

Such changes include:

- Reduction of required details on existing production infrastructure usage, extensive code scaling examples, emphasis on scientific excellence and publications using such codes.
- Inclusion of more details on porting and optimization of applications with the inclusion of questions such as:
  - Description of main algorithms used and details on their implementation
  - Description of the known or potential performance bottlenecks
  - Description of possible solutions for performance issues
  - Description of I/O strategies used by the codes.
  - Description of required support from PRACE.
  - Description of the potential technology transfer in future.

Although as described above several steps toward the realization of preparatory access to Tier-1 resources have been performed, the actual implementation has not been decided yet and the subject has been left open for further evaluation, after parameters consultation with the users takes place in order to identify the real needs for the introduction of such an access method for Tier-1 resource.