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 Project

Grant Agreement Number: RI-283493

D2.1
Migration from DEISA2 to PRACE-2IP

Final

Version: 1.0
Author(s): Juha Fagerholm, CSC, Alison Kennedy, EPCC
Date: 21.12.2011

Project and Deliverable Information Sheet

PRACE Project	**Project Ref. №: RI-283493**	
	Project Title: PRACE Second Implementation Project	
	Project Web Site: http://www.prace-project.eu	
	Deliverable ID: < **D2.1**>	
	Deliverable Nature: < Report >	
	Deliverable Level: PU *	**Contractual Date of Delivery:** 31 / 12 / 2011
		Actual Date of Delivery: DD / Month / YYYY
	EC Project Officer: Bernhard Fabianek	

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

Document Control Sheet

Document	**Title: Migration from DEISA2 to PRACE-2IP**	
	ID: D2.1	
	Version: <1.0 >	**Status:** ***Final***
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2007	
	File(s): D2.1.docx	
Authorship	**Written by:**	Juha Fagerholm, CSC, Alison Kennedy, EPCC
	Contributors:	Lilit Axner, KTH-PDC; Anni Jakobsson, CSC; Chris Johnson, EPCC; Petri Nikunen, CSC; Wim Rijks, SARA
	Reviewed by:	Florian Berberich, JSC Jorge Naranjo, BSC
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	02/December/2011	Draft	Some parts still missing
0.2	03/December/2011	Draft	
0.3	05/December/2011	Draft	
0.4	06/December/2001	Draft	Additions from Alison
0.5	07/December/2011	Draft	Addition by Chris
0.6	08/December/2011	Draft	For internal review
1.0	21/December/2011	Final version	

Document Keywords

Keywords:	PRACE, HPC, Research Infrastructure

Disclaimer

This deliverable has been prepared by Work Package 2 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

Table of Contents

List of Figures

List of Tables

References and Applicable Documents

[1] PRACE-1IP deliverable D4.3.1 Cross National Programme for Access to Tier-1 Resources
[2] www.prace-ri.eu
[3] PRACE-2IP Grant agreement Annex I – "Description of Work"

List of Acronyms and Abbreviations

AISBL	Associations Internationales Sans But Lucratif (International association without lucrative purpose)
API	Application Programming Interface
BSC	Barcelona Supercomputing Center (Spain)
BSCW	Be Smart – Cooperate Worldwide (Public Cooperation Platform)
CEA	Commissariat à l'Energie Atomique (represented in PRACE by GENCI, France)
CINECA	Consorzio Interuniversitario, the largest Italian computing centre (Italy)
CINES	Centre Informatique National de l'Enseignement Supérieur (represented in PRACE by GENCI, France)
CPU	Central Processing Unit
CSC	Finnish IT Centre for Science (Finland)
CSCS	The Swiss National Supercomputing Centre (represented in PRACE by ETHZ, Switzerland)
DPMDB	DECI Project Management Database
DECI	Distributed European Computing Initiative
DEISA	Distributed European Infrastructure for Supercomputing Applications. EU project by leading national HPC centres
DoW	Description of Work
EC	European Community
EPCC	Edinburgh Parallel Computing Centre (represented in PRACE by EPSRC, United Kingdom)
EPSRC	The Engineering and Physical Sciences Research Council (United Kingdom)
ETHZ	Eidgenössische Technische Hochschule Zürich, ETH Zurich (Switzerland)

FZJ	Forschungszentrum Jülich (Germany)
GCS	Gauss Centre for Supercomputing (Germany)
GENCI	Grand Equipement National de Calcul Intensif (France)
GPFS	General Parallel File System (by IBM)
HLRS	High Performance Computing Center Stuttgart (represented in PRACE by GCS, Germany)
HP	Hewlett-Packard
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing
IBM	Formerly known as International Business Machines
ICHEC	Irish Centre for High-End Computing (represented in PRACE by NUI Galway, Ireland)
IDRIS	Institut du Développement et des Ressources en Informatique Scientifique (represented in PRACE by GENCI, France)
JSC	Jülich Supercomputing Centre (FZJ, Germany)
KTH	Kungliga Tekniska Högskolan (represented in PRACE by SNIC, Sweden)
LINPACK	Software library for Linear Algebra
LRZ	Leibniz Supercomputing Centre (Garching, represented in PRACE by GCS, Germany)
NCF	Netherlands Computing Facilities (Netherlands)
NUI	National University of Ireland – Galway
PDC	Parallelldatorcentrum (KTH, represented in PRACE by SNIC, Sweden)
PDF	Portable Document Format
PHP	PHP: Hypertext Preprocessor
PI	Principal Investigator
PMO	Project Management Office
PRACE	Partnership for Advanced Computing in Europe; Project Acronym
PSNC	Poznan Supercomputing and Networking Centre (Poland)
RTF	Rich Text Format
RZG	Rechenzentrum Garching, Max-Planck-Gesellschaft (represented in PRACE by GCS, Germany)
SARA	Stichting Academisch Rekencentrum Amsterdam (Netherlands)
SE	Scientific Evaluation
SGI	Silicon Graphics, Inc.
SNIC	Swedish National Infrastructure for Computing (Sweden)
STFC	Science and Technology Facilities Council (represented in PRACE by EPSRC, United Kingdom)
SUSP	Scientific Users' Selection Panel (in HPC-Europa)
TE	Technical Evaluation
Tier-0	Denotes the apex of a conceptual pyramid of HPC systems. In this context the Supercomputing Research Infrastructure would host the Tier-0 systems; national or topical HPC centres would constitute Tier-1

Executive Summary

The DECI programme for high-end scientific projects in HPC was started in 2005 by the EC-funded DEISA project. The programme was based on the access to the top-level European HPC systems. Within the DEISA project, the programme had six calls which proved to be very successful. After the DEISA projects ended in April 2011, the DECI programme has been migrated to PRACE. The first PRACE DECI pilot call, DECI-7, was opened already in May 2011 by PRACE-1IP synchronously with the PRACE 3rd regular call. In total, 54 proposals were received. The Principal Investigators of these proposals represented 17 European countries. The call was oversubscribed by a factor over 2.5. Of the proposals, 35 DECI-7 projects were selected, beginning to run on November 1st, 2011, on the Tier-1 systems provided by 15 PRACE partners who contributed over 90 million CPU core hours.

This deliverable describes also the processes involved in the DECI project management as they have been migrated to PRACE. The PRACE operational environment on the Tier-1 systems is not described.

The first pilot call was prepared in a very short time. Experiences from the pilot call were used by for the launching of the next, DECI-8, call in November 2011. The processes for the launch need still to be refined, and improvements have been planned for the forthcoming call in 2012. The proposal submission was based on email and MS Word proposal templates as in DEISA due to lack of preparation time. The templates were modified from the DEISA templates to PRACE type templates. If the amount of proposals grows in the future, the collecting of forms by email will be too tedious and a more advanced database system is needed, such as what is used for PRACE Tier-0 proposals or HPC-Europa proposals.

The technical evaluation or review was performed by experts in the home sites of the proposals. The notion of a home site is specific to DECI and means a site in the country of affiliation of the Principal Investigator. For other countries, a list of home sites is used. The technical evaluation form was adapted from the one used in DEISA. After the experiences in DECI-7, improvements have been introduced for the next call. The scientific evaluation or peer review is different from PRACE Tier-0 calls. The proposals with a PI from DECI countries were evaluated by a national peer review committee countries except proposals from Poland; proposals from countries outside the DECI countries and Poland were reviewed by HPC-Europa SUSP.

For computer resources, a commitment by the DECI partners is needed. The resources are normalised by a conversion factor which is based on processor speed and the overall architecture. This is adapted also from DEISA but new factors have been agreed within PRACE. In DECI-7, 15% of the resources were reserved for projects external to DECI countries. Rest of the resources were exchanged between DECI countries. The best projects were selected based on the ranking by the scientific evaluation. Projects were entered in the DECI Project Management database set up at SARA; it is used to follow their progress.

Application support and porting is provided by the PRACE-2IP project WP7 task T7.2. This support starts with the proposal phase and continues for accepted projects in meetings with the project group to present the services of PRACE. Special support is given to projects that have asked for application enabling. All projects will be followed in monthly video conferences where both WP2 and T7.2 are present. Project reports will be collected and stored in the BSCW for PRACE use.

Discussion on future improvements has been included on proposal submission and launching the calls. Improvements are also related to the growing number of partners and proposals and

collaboration with the Tier-0 project management. To clarify the future of DECI, the possibility of the PRACE Optional Programme is recommended to be investigated by WP2.

1 Introduction

The purpose of the deliverable is to describe the transition of the single-project access programme known as DECI from the DEISA project to the PRACE-2IP project with a special focus on the processes and procedures. The document is intended for three audiences. Firstly for the project management board, which has responsibility for running the calls and for selecting and supporting a target number of DECI projects, secondly for the PRACE AISBL who have responsibility for managing the resources and calls of the European HPC infrastructure in a cohesive and integrated manner and thirdly for the project personnel who have responsibility for devising and managing the DECI processes.

The initial priority for PRACE was to integrate DECI into PRACE so as to maintain the momentum of the programme and provide continuity of access for the European research community. This integration of DECI into PRACE was undertaken as a pilot, with the existing DECI processes (as defined during the DEISA projects) adapted minimally to fit the PRACE project structure. The DECI pilot call was confirmed by the PRACE Council and opened simultaneously with the Tier-0 call.

In this deliverable, produced in the third project month of the PRACE-2IP project, we will review and assess the longer-term suitability of these processes and make recommendations as to how they can be improved so as to better position DECI for its future role in the HPC ecosystem. By producing an early review, recommendations and suggestions can quickly be acted upon for future DECIs.

The structure of the document is as follows. In Section 2, first the DECI history and the meaning of DECI in the PRACE context is described, thereafter a general description of the work for DECI and results of the DECI-7 call are presented. In the Section 3, the procedures for DECI are described in detail. In Section 4, how to improve the processes is discussed. Section 5 gives the conclusions.

DECI is performing strongly and is well regarded and supported by the scientific community in HPC. The challenge for PRACE is to take the legacy of DEISA and to further develop DECI into an instrument of PRACE with a well-understood role in the HPC ecosystem. This deliverable recommends that an Optional Programme of PRACE AISBL be considered as a means of defining the role and scope of DECI in the future and in giving the Tier-1 sites a mandate to deliver a programme which meets these objectives.

2 Role of DECI

2.1 Brief history of DECI

DECI, originally an abbreviation for the DEISA Extreme Computing Initiative, was launched under the DEISA EU project in 2005 as the first European programme in high performance computing for complex, demanding and innovative simulations. This highly successful initiative was based on the fastest supercomputers in Europe, and paved the way to the PRACE project and its regular project calls. The DEISA project, including DEISA, eDEISA, and DEISA2 EU projects, opened six DECI calls; the subsequent DECI projects were in production in 2005–2011. At the time of its inception, DECI offered access to the top level

HPC systems in Europe – the machines which were made available through DECI during DEISA were almost exclusively "Top 100" machines.

In PRACE, the role of DECI, now an abbreviation for the Distributed European Computing Initiative, is to provide Tier-1 resources as PRACE project access will provide the top Tier-0 high-performance computing resources. The definitions of a Tier-1 centre and Tier-1 resources are given by PRACE-1IP WP4 deliverable D4.3.1 [1]. While the larger European Tier-1 systems occupy high positions in the "Top 100", DECI resources have been augmented both by a number of "older" HPC systems (i.e. systems which were previously the most powerful national resources but which have largely been supplanted or augmented for national use by newer systems) and by a number of national HPC systems, typically clusters, from smaller PRACE countries.

DECI's original goal was to enable European computational scientists to obtain access to the most powerful national computing resources in Europe regardless of their country of origin or work. This objective has been evolving over the past three years since the introduction of the PRACE Project Access programme, which provides access to Tier-0, although there are some differences between the two programmes.

DECI has therefore a new objective or series of objectives. One important part of its remit is to offer a ramp to Tier-0, by providing access to medium to large HPC systems of similar architectures to the PRACE Tier-0 systems. A second objective is to provide access to state-of-the-art HPC systems to computational scientists who do not require either the volume of CPU cycles typically awarded by Tier-0 or access to as large a number of cores to undertake world-leading science. A third objective is to foster collaborative science by providing access to consortia of scientists from different countries. The consortia who apply to DECI are often new collaborations whereas the ones who apply to Tier-0 are typically much larger and better established. A fourth objective is to offer access to a wider range of systems architectures and machine configurations than can be offered at the Tier-0 level.

DECI selects a number of capability computing projects by scientific peer-review on the basis of innovation and scientific excellence. DECI's ability to provide application support and training to these selected projects, as part of a package of support to selected projects, has been one of the main assets in its six years of successful history in DEISA.

2.2 From DEISA to PRACE DECI

DECI-7 was intended as a "Pilot call" and enabled a transition from DECI within the DEISA to DECI within PRACE. This meant that it was possible to reuse much of the DEISA infrastructure, both in terms of the DECI process and for the physical infrastructure itself. However, there were also a number of challenges as many of the partners were new to DECI and even within ex-DEISA sites some of the individual staff were new to DECI. In addition, although some of the supporting technology was re-used, some of it was moved from one site to another. For example, the BSCW was moved from CSC to SARA and the DPMDB was moved from RZG to SARA.

Integrating new partners and staff into DECI was a gradual process. Some partners had been partly involved under DECI-6 by providing compute resources but having no projects as home site and not being fully part of the DEISA infrastructure. Most of these partners had taken part in the monthly video-conferences during the final stages of the DEISA project and so had already been exposed to the DECI process. Other partners were completely new and so the process has had to be learnt more quickly. To aid the integration of new partners and also to smooth over the transition from DEISA to PRACE, we continued monthly video

conferences throughout the interim period between DEISA and PRACE and invited all the new partners along.

The process of selecting proposals was based on technical and scientific evaluation, both of which are described elsewhere in this document. The procedures used were based on the equivalent procedures within DEISA. The main differences were due to the increased scale of the project in terms of the number of partners and machines involved.

Once the technical and scientific evaluations had all been completed, the next task was to determine how many projects could be accepted and what resources should be assigned to each project both qualitatively in terms of assigning machines and quantitatively in terms of the amount of CPU-hours to provide. Each internal proposal was listed according to the scientific rankings which themselves had taken into account the technical evaluations. In order to decide which projects to award time to, two guiding principles were used. We considered that both scientific excellence and the notion of *juste retour* were the most important factors to consider. To achieve this, each centre considered its own list together with the amount of resources that the centre was contributing to internal DECI projects. From this information it was possible for each centre to decide how far down its ranked list of proposals the projects could be accepted. By considering the comments in the technical and scientific evaluations, in some cases certain projects had the amount of resources cut back in order to accommodate more projects. This was generally done towards the lower end of the accepted projects' lists in order that the best projects should receive their full request where possible. For external projects the same principle was applied with each centre that was responsible for a particular project providing information on the extent to which the project's resources could be scaled back.

This procedure generally worked well but the change in the number of contributing partners and the wide discrepancies in the amount of CPU contributed did exacerbate some problems.

In some cases countries and centres were contributing large amounts of computing resource but only had a small number of projects. This was a situation which never arose in DEISA, where the amount of CPU requested by applicants from a country always exceeded the amount of resources contributed by that country. However, in DEISA, where the centres were the partners, it was clear where the responsibility for publicity lay. In DEISA, the External Relations work package undertook broadcast marketing and each centre undertook targeted marketing at a national level. This division of responsibilities is not so clear in PRACE where the PRACE partner often has sub-contracted partners, so perhaps DECI was not marketed at a national level as effectively as it could be. However, we have to consider two possibilities – firstly, that there are some countries where there is very low demand for cross-national Tier-1 resources or secondly, that there is a demand but that the researchers are not fully aware of the opportunities. We need to understand the situation better by working with the countries concerned in order to respond in an appropriate way, either via training, applications support, better publicity etc.

To deal with mismatch between supply and demand at a national level, some flexibility was required and we agreed that the remaining CPU time should be used to assist centres with a large number of projects. Any imbalance caused by this action for an individual DECI can be redressed at future DECIs. It was thus implicitly recognised that the notion of *juste retour* should not be applied on a DECI by DECI basis but should span a longer agreed time period.

Once the amount of resources for each accepted project had been agreed we were in a position to assign projects to machines. There are a number of constraints on this assignment problem as each project has its own architecture requirements and fitting 35 projects across the 19 available machines becomes an over-constrained problem with no perfect solution. To help with this problem, for this DECI we introduced a new step which was for each centre to come

up with a ranking for the classes of architecture which would best suit each accepted project. From this it was possible to come up with a reasonable assignment of machines which attempted to match projects to resources without splitting projects over machines to too great an extent. Other factors which needed to be taken into account were storage requirements, package availability, maximum job length and number of CPUs required which sometimes tied projects to specific machines. Even after this process, once enabling work or production runs start it will occasionally turn out that some machines are not suitable and resources will have to be exchanged between projects with the agreement of all parties involved. This can happen at any time, even well into the production phase of the project. This situation previously occurred in DEISA, however it will be made more complicated by the larger number of systems and by the small number of projects running on many of the systems.

In order to exchange resources between countries, it is necessary to have an agreed exchange rate mechanism. For this pilot call, we used the exchange rate mechanism established by DEISA, since the majority of machines in DECI-7 either had been in DEISA (and so had an agreed exchange rate) or were of an identical architecture to another machine contributing resources to DECI-7. However, relying on historical information from DEISA is not a sustainable position for PRACE as new machines continue to be added and the issue of setting exchange rates needs to be re-visited. (DEISA used a combination of Linpack ratings, analysis of processor speeds and the geometric means of the results of running the DEISA benchmark suite on each new machine as it was integrated into the infrastructure to come up with a conversion factor.)

At all stages, the projects accepted and machines assigned were agreed amongst the partners by email and video conference and the decisions were then ratified by the PMO.

At several stages in the process, information has to be entered into the DPMDB by each centre. This happens in at least three stages. Initially when proposals are first received basic information is added about each proposal (PI, links to proposal, etc). Later when the list of accepted projects is decided this information is entered. Finally, when machines are assigned this information is also entered. With many participating centres, several of them new to DECI, this procedure can take quite a long time to complete. In the future we may have to consider the balance between information entered locally by each centre and information entered centrally by the DPMDB administrator or DECI management. It should be noted that there will new partners at least until DECI-9.

2.3 Summary of the DECI pilot, or DECI-7, call

The DECI pilot, or DECI-7, call was the first synchronized PRACE regular and DECI call after the DEISA project concluded in April 2011. The call was published on 2 May, 2011, leaving very little time to coordinate the first synchronized call.

Despite the time pressure in the call launch process, 54 proposals were received in the DECI-7 call, and 35 projects were awarded 91 million processor core hours in total. Proposals were from a wide variety of disciplines, see Figure 1: Proposals by scientific discipline in DECI-7. Principal Investigators (PIs) of the proposals represented 17 countries (Table 1: PIs by country). The call was oversubscribed by more than a factor of 2.5. Due to the excellent quality of the proposals, the contributing partners increased the resources offered to the call by 11% over the original commitment; however this still meant that many good projects could not be granted resources.

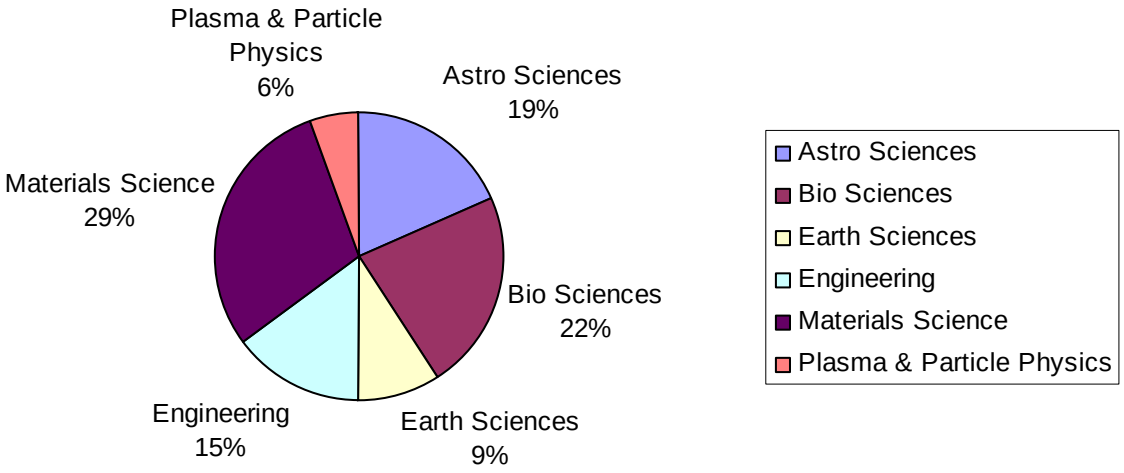

Figure 1: Proposals by scientific discipline in DECI-7

Finally, 35 DECI-7 projects were awarded Tier-1 resources. Of these, 30 projects were internal, i.e. from those countries that contributed to DECI-7 resources, and 5 projects were external to DECI-7 countries. The computational resources (core-hours) awarded to projects are shown in Table 2: Time awarded to projects with their respective home and execution sites.

Country of PI	Number of PIs
Austria	1
Belgium	2
Cyprus	1
Finland	3
France	3
Germany	8
Greece	1
Hungary	1
Ireland	2
Italy	11
Poland	1
Portugal	3
Spain	2
Sweden	3
Switzerland	2
The Netherlands	4
UK	6
Total	54

Table 1: PIs by country

Internal/External	DECI project	Computational resources awarded (core-hours)	DECI home site	DECI execution site(s)
External	PICKH	1,500,000	IDRIS	CINECA
External	DIAVIB	1,080,000	HLRS	CINECA,SARA
External	VIRonSAMs	4,480,000	HLRS	CINES,CINECA
External	BlackHoles	4,200,000	BSC	PSNC
External	PHOTMAT	5,062,500	HLRS	PSNC,CINES
Internal	CatDesign	1,292,000	BSC	LRZ,PSNC SGI
Internal	MIXTUDI	3,750,000	CINECA	FZJ,CINECA
Internal	MAESTRO	2,000,000	CINECA	IDRIS
Internal	SCW	1,410,000	CINECA	LRZ
Internal	PETAHUB	1,800,000	CINECA	NCSA,PDC
Internal	NUWCLAY	1,500,000	CINES	CINECA
Internal	ElmerIce	1,400,000	CINES	PDC
Internal	WESF	1,200,000	IDRIS	RZG,CSC
Internal	Planck-LFI	3,500,000	CSC	CSC
Internal	TanGrin	3,500,000	CSC	HECToR
Internal	EC4aPDEs-2	2,500,000	EPCC	BSC,CSC,HLRS,IDRIS
Internal	HYDROGEN-ILs	1,314,000	EPCC	PDC
Internal	HELIXKINETICS	2,713,190	EPCC	ICHEC,FZJ
Internal	HIFLY	2,100,000	EPCC	PSNC
Internal	HIGHQ2FF	5,000,000	FZJ	PDC
Internal	NR-NSNS-BHNS	3,000,000	FZJ	CSC
Internal	LGICTAMD	3,024,000	ICHEC	ICHEC
Internal	NANOBIO-2	2,000,000	ICHEC	HLRS,CINES
Internal	CASiMIR	3,263,148	LRZ	SARA
Internal	DiSMuN	3,750,000	PDC	PDC,SARA
Internal	SIVE-2	6,250,000	PDC	HECToR
Internal	MUSIC	231,000	PDC	IDRIS
Internal	SPIESM	3,750,000	PDC	PDC
Internal	SIMONA	600,000	PSNC	SARA,PSBC,IDRIS
Internal	ARTHUS-3	2,750,000	RZG	FZJ
Internal	EUTERPE-4	1,500,000	RZG	FZJ
Internal	SMARC	700,000	RZG	LRZ
Internal	LASIPROD	700,000	RZG	LRZ
Internal	RBflow-2	6,000,000	SARA	RZG,SARA
Internal	HRPIPE	1,800,000	SARA	PDC
	Total	90,619,838		

Table 2: Time awarded to projects with their respective home and execution sites. Each project has a home site which is typically a PRACE site in the country where the PI is affiliated.

3 Procedures of the DECI calls

The DECI pilot call was launched by PRACE-1IP WP4 but the follow-up was taken by PRACE-2IP WP2 as the 2IP project started in September 2011.

In line with DEISA DECI (or existing PRACE Tier-0) practice, the following procedures were implemented:

- Every proposal underwent a technical evaluation (TE), performed by an HPC expert from one of the PRACE partners;

- Every proposal underwent a scientific evaluation (SE), performed by a peer review committee recognised by PRACE;
- Every PRACE partner with a Tier-1 system explicitly opted in or out of the pilot call and those who opted in pledged a fraction of their total resources to the call, for use between 1 November 2011 and 31 October 2012. These resources are initially expressed in local core-hours.
- CPU resources were classified into four architecture classes – IBM BG/P, Cray XT/XE, IBM P6 and Clusters.

In the chapters below, a more detailed description on the processes and actions during DECI-7 is outlined. The actual work in DECI-7 has been divided over several work packages in PRACE-1IP and PRACE-21P but involves also PRACE AISBL and the DECI sites. Since PRACE-2IP started in September 2011, the coordination of DECI calls and projects is provided by 2IP WP2.

3.1 Call launch and publicity

3.1.1 *DECI pilot or DECI-7 call*

The call text of the DECI pilot call was prepared mainly by PRACE Tier-0 representatives in PRACE AISBL, with input from WP4 of PRACE-1IP. The call information approved by PRACE AISBL was published on May 2nd. The call text included a press release, published on the PRACE RI website www.prace-ri.eu [2] by the PRACE press team (prace-press-team@fz-juelich.de), and a PDF file containing more detailed information about the call. The call was also disseminated to PRACE's already existing press contacts, the AlphaGalileo service and disseminated via PRACE partner sites' local dissemination channels. Researchers who have subscribed to "PRACE Call Alerts" at the PRACE website also received information about the new call. Language translations were done by some of the PRACE partners in order to disseminate the call locally. The PRACE web team (prace-web-team@fz-juelich.de) prepared and published a banner to the website to advertise the call. There was therefore an integrated call text for Project Access (Tier-0) and for DECI for the first time.

There were afterwards requests from Tier-1 (i.e. PRACE-1IP WP4 and DECI sites) to do changes to the text on the PRACE website regarding information related to Tier-1. However, the official PRACE approval process at the time (PRACE AISBL Board of Directors, BoD) never approved the changes, thus the changes could not be applied to the text on the website. This made it difficult to clarify some aspects of the call which were not very obvious to respondents, in particular to provide clear guidance on whether researchers should apply to Tier-1, Tier-0 or could apply to both. Information about the PRACE 3rd call results was published by the PRACE press team.

3.1.2 *DECI-8 call*

Due to the problems of the tight schedule in the DECI-7/PRACE 3rd regular call, it was decided that the call process needs to be clarified, taking both Tier-0 and Tier-1 representatives into account. At the Barcelona 1IP all-hands & 2IP kick-off meeting in September 2011 decisions on how to improve the process for the DECI-8/PRACE 4th regular call were proposed by PRACE-2IP WP3 and approved. Two persons from Tier-0 side and two persons from Tier-1 side were named to prepare the text for the call documents. PRACE press and web teams were involved in the process to publish the call on time via the existing dissemination channels.

A schedule for the call process was sent by the press team to all parties and the schedule was agreed after some minor date changes later approved:

- Call draft ready by October 10th
- Comments submitted and final version ready by October 14th
- PRACE AISBL's approval by October 21st
- October 23rd–30th, time for language translations by PRACE partners & preparations by the web team.
- Call published on November 2nd.

However, there were delays in the schedule due to various reasons. These involved the necessary discussions between PRACE AISBL and the PRACE project on the role of Tier-0 and Tier-1 parts in the call text which took more time than anticipated. Also the necessary decisions by the PRACE project management board on synchronizing the DECI-8 and PRACE 4th regular calls and on launching the calls had to be made before the call. The issues caused by the delays have been recognised and discussions have been held within the project for future calls.

3.2 Proposal templates

Although the DECI-7 call was combined with the PRACE 3rd regular call, the gathering of proposals and their evaluation are two separate processes. Unfortunately due to the very short preparation time, and the lack of funded project effort for this task, it was not possible to adapt the electronic application form that is used for the Tier-0 calls to the specific needs of the Tier-1 call. Therefore the application form for the DECI-7 call was still based on the ones used for the DEISA DECI calls, which were in MS Word format, but adapted for this specific call. This implied that the applicants still had to fill in the MS Word form, and sent it in by e-mail as a MS Word or PDF document. DECI staff had then to gather the proposals by hand and store them in the PRACE-2IP BSCW.

While processing the applications it became apparent that there still were some slight flaws and inconsistencies in the layout of the form, so all the comments from DECI staff were collected and the form for the DECI-8 call was adapted accordingly (see 6.1).

Evaluation of the whole process of collecting forms in Word format led to the conclusion that this is a tedious process that is far from ideal and introduces numerous opportunities for human error as information provided by applicants needs to be cut, pasted and reformatted before it can be stored in the project database. Therefore we are looking into the possibilities to switch to electronic application forms, possibly integrating the application form for the DECI proposals with the one that is used for Tier-0 applications.

3.3 Technical review

The technical review of DECI proposals is done during the two week period, after the closing date of the corresponding DECI call. The task is assigned to the WP7 task T7.2, subtask (A) of the PRACE-2IP project. Representatives from each DECI site as well as other PRACE HPC centres that have expressed their wish to be active in WP7 task T7.2 are identified for the technical review subtask.

Technical evaluation (TE) is currently based on a Word document form (see 6.2) that was used in the DEISA DECI process, but adapted to the altered layout of the DECI-7 proposal form. All the fields within this form have been discussed, fine tuned and finally defined by

WP2. After the process of technical evaluations of the DECI-7 proposals was finished a number of improvements have been introduced for the DECI-8 call, based on the experience of the evaluators. Each site representative in the T7.2 (A) uses this form for evaluation of a given DECI project application. The completed TEs are being uploaded to the PRACE BSCW by T7.2 (A) members. These TEs together with the project proposals are sent to the peer review committees (as defined in the section 3.6 Peer Review) by WP2 representatives. Parallel to this T7.2 (A) staff has populated the DPMDB database with the applications' technical information to be used in future by WP2 and WP6.

Considering the continuing nature of technical review task, WP2 together with T7.2 (A) is in discussion about enhancement and simplification of the procedure. One of the simplifications in discussion is the implementation of a web based technical review form. This has already been practised in the projects like HPC-Europa. Content-wise, the web based TE form will mirror the fields defined in the current Word document version of the TE form, but it will greatly simplify the process of collecting, storing and maintaining all kind of documents in different places.

Current practice is to divide the DECI proposals for Technical Evaluation based on the home site of the proposals. Because the allocation of home sites to different DECI partners is rather inhomogeneous, this 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 with the preferred execution platform and form pools of technical experts for each architecture that will do the evaluation for these applications.

Another point to consider is the current discussion within a PRACE Working Group organised by Axel Berg of WP6, about making the process of application for PRACE resources more transparent for users. Possible options for this are being discussed. This can have a direct impact on the current technical review procedure.

The next upcoming technical evaluation will be conducted for DECI-8 call during the last two weeks of January as the closing date of DECI-8 call is 10th January.

3.4 Peer review

All DECI proposals are subject to a technical and a scientific evaluation. To divide the workload the DECI-7 proposals, just as all previous DEISA DECI proposals, were allocated to one of the partners in DECI, who then acts as home site for that project. The home sites (ICHEC, EPCC, BSC, RZG, FZJ, HLRS, LRZ, IDRIS, CINES, PSNC, CINECA, CSC, SNIC, SARA) were responsible for appointing a peer review committee that would perform the scientific evaluation of all the projects allocated to them, which in most cases coincided with all proposals from a particular country, the exceptions being France and Germany. The proposals of the two French sites (CINES, IDRIS) were combined and evaluated by a national evaluation committee, while the proposals of the four German sites (FZJ, RZG, HLRS and LRZ) were evaluated and ranked separately for each site.

The Scientific Peer Review Committee could be either a National Committee or sites could delegate the process to HPC-Europa SUSP. The evaluation of the external (non-partner) projects was done by default by HPC-Europa SUSP. All sites, except Poland, decided to let the scientific evaluations be performed by a National Committee.

Unlike in the DEISA DECI process, where the technical and scientific evaluations were done in parallel, in DECI-7 we chose to perform the technical evaluations first, (as is done for Tier-0) giving the Scientific Peer Review Committees the opportunity to use the technical evaluation results as additional input. The output of each Scientific Peer Review Committee

was a review report for each individual proposal they had to evaluate, and most importantly for the DECI process, a ranking of all the proposals that they had reviewed.

The collection of external projects was reviewed separately by HPC-Europa SUSP and also received a separate ranking.

The ranking of all proposals was used as basis for the process of accepting or rejecting the proposals.

3.5 Resource contributions

In DECI, the contribution of any PRACE partner to each DECI call has to be confirmed by the management of the PRACE project. As a limitation the DoW [3] specifies that sites providing Tier-1 resources for DECI should provide enough HPC capacity to accommodate at least one DECI project per call. However, a partner may participate e.g. to only every second call.

In DECI-7, over 90 million CPU core hours were contributed by 15 partners for PRACE. Contributions per partner varied from 1.25% to 5% of their total Tier-1 CPU resources (see Table 3: DECI-7 computer resources).

System		Arch. Class	no of cores	% annual comm.	%DECI-7 committed	core-hours
Juropa	Bull Nehalem Cluster	Cluster	17664	5 %	2,50 %	2 900 000
vip	Power6	Power 6		5 %	2,50 %	1 150 000
genius	BG/P	BG/P		5 %	2,50 %	2 872 000
	SuperMUC Migration System	Cluster	8200	5 %	3,00 %	1 437 000
Laki	NEC Nehalem Cluster	Cluster		5 %	2,50 %	784 000
BABEL	BG/P	BG/P	40960	5 %	2,5%	7 176 192
JADE	SGI ICE EX8200	Cluster	23040	2,50 %		2 000 000
JADE-1		Cluster	12288	2,50 %	1,25 %	1 000 000
JADE-2		Cluster	10752	2,50 %	1,25 %	1 000 000
HECToR	Cray XE6	Cray		5 %	2,50 %	7 800 000
Mare Nostrum	Cluster	Cluster		5 %	2,50 %	1 900 000
Louhi	Cray XT4/5	Cray	10864	5 %	3,00 %	2 284 000
Huygens	IBM p575 HydroCluster	Power 6		8 %	4,00 %	880 000
Lindgren	Cray XE6	Cray	36384	5 %	5,00 %	12 749 000
	IBM SP6	Power 6		8 %	4,00 %	1 400 000
	IBM Hybrid Westmere Cluster	Cluster		3 %	4,00 %	850 000

System		Arch. Class	no of cores	% annual comm.	%DECI-7 committed	core-hours
	SGI UV1000	Cluster		10 %		2 747 000
	HP	Cluster		15 %		1 513 728
Stokes	SGI ICE EX8200	Cluster	380	5 %		3 363 840
EA"ECNIS"	BG/P	BG/P		5 %		2 870 000

Table 3: DECI-7 computer resources

In addition to this, large amounts of storage space and application support were provided, but they were counted as supplementary resources to computing time. The total amount of application support is the same as the person months in T7.2. Contributed resources were entered in the DECI Project Management Database (DPMDB).

In addition, conversion factors, based on processor speeds and the overall architecture, have also been agreed. Where a conversion factor already existed for a processor in DEISA DECI, this same factor has been carried forward to PRACE DECI. Factors for new processors have been discussed and agreed within PRACE. The purpose of the conversion factor is to facilitate exchange of CPU resources between partners by means of establishing a common PRACE CPU currency. This unit will additionally be used for internal accounting and reporting. These normalised units are based on performance relative to the IBM Power4 (P4) processor for historical reasons.

Of the 21 partners of PRACE, 15 participated in DECI-7. Further new partners will join in the coming DECI-8.

Note that allocations and commitments by country will never match exactly, since 15% of the resources committed were set aside for projects from non-contributing countries (see Figure 2: DECI-7 commitments by country and Figure 3: DECI-7 allocations by country).

3.6 Resource exchange and allocations

In DECI-7, 15% of computing resources were reserved for external projects, projects that come from countries that do not participate in DECI. The rest of the resources were exchanged between countries (*juste retour*).

Projects were assigned to machines according to their suitability to computer architectures and the amount of resources available on machines. Availability of storage resources also played a role also in the assignment process.

After the projects were assigned, they were entered in the DECI Project Management Database (see Figure 4: DPMDB view on DECI-7 projects).

During DECI's within DEISA, it turned out occasionally that resources given to projects were not suitable or that resources were unavailable for unforeseen reasons. In DECI-7 we have not run into that problem, but the projects are only just starting. If such an incident occurs, a new execution site should be arranged for the project. Probably two or more projects need to swap execution sites, or some execution site offers additional resources to host the project.

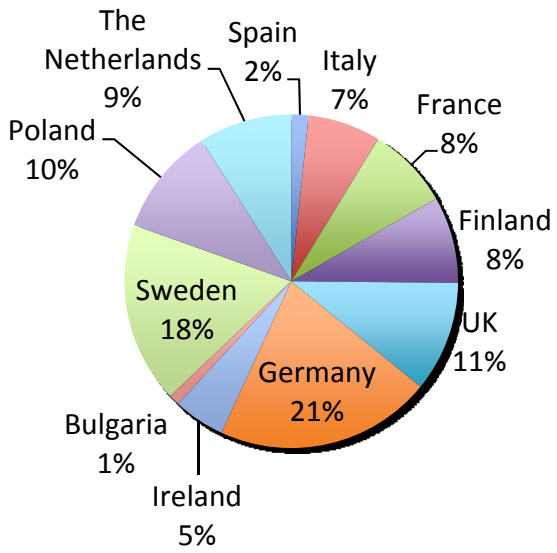

Figure 2: DECI-7 commitments by country

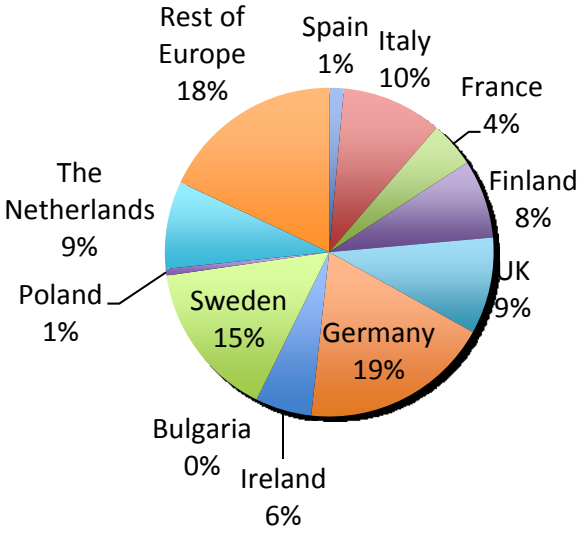

Figure 3: DECI-7 allocations by country

DPM-Database Welcome, Juha Fagerholm Database: DECI-7 (2011) access / admin access

Projects Survey DECI-7 (2011)

(Full survey, sorted by projects, standardized Core hours)

Int. #	Projectname	DEISA Site (home)	DECI granted Core hours	consumed Core hours	consumed until Jul 11	DEISA Site (exec)	assigned Core hours	consumed Core hours	consumed until Jul 11	Project state
90469	ARTHUS-3	RZG	2,750,000	0	~ 0 %	FZJ	2,750,000	0	~ 0 %	exec assigned
90470	BlackHoles	BSC	4,200,000	0	~ 0 %	PSNC	4,200,000	0	~ 0 %	*not yet started*
90471	CASIMIR	LRZ	3,263,148	0	~ 0 %	SARA	3,263,148	0	~ 0 %	*not yet started*
90472	CatDesign	BSC	1,292,181	0	~ 0 %	LRZ	775,200	0	~ 0 %	*not yet started*
						PSNC	516,800	0	~ 0 %	*not yet started*
90475	DIAVIB	HLRS	1,080,000	0	~ 0 %	CINECA	648,000	0	~ 0 %	exec assigned
						SARA	432,000	0	~ 0 %	exec assigned
90476	DISMuN	PDC	3,750,000	0	~ 0 %	SARA	1,578,000	0	~ 0 %	In setup
						PDC	2,172,000	0	~ 0 %	In setup
90478	EC4aPDEs-2	EPCC	2,500,000	0	~ 0 %	BSC	1,520,000	0	~ 0 %	exec assigned
						CSC	665,000	0	~ 0 %	exec assigned
						HLRS	235,000	0	~ 0 %	exec assigned
						IDRIS	80,000	0	~ 0 %	exec assigned
90480	Elmerice	CINES	1,400,000	0	~ 0 %	CSC	1,400,000	0	~ 0 %	*not yet started*
90481	EUTERPE-4	RZG	1,500,000	0	~ 0 %	FZJ	1,500,000	0	~ 0 %	exec assigned
90484	HELIXKINETICS	EPCC	2,713,190	0	~ 0 %	FZJ	1,199,230	0	~ 0 %	exec assigned
						ICHEC	1,513,960	0	~ 0 %	exec assigned
90486	HIFLY	EPCC	2,100,000	0	~ 0 %	PSNC	2,100,000	0	~ 0 %	exec assigned
90487	HIGHQ2FF	FZJ	5,000,000	0	~ 0 %	PDC	5,000,000	0	~ 0 %	exec assigned
90489	HRPIPE	SARA	1,800,000	0	~ 0 %	PDC	1,800,000	0	~ 0 %	exec assigned
90490	HYDROGEN-ILs	EPCC	1,314,000	0	~ 0 %	PDC	1,314,000	0	~ 0 %	exec assigned
90491	LASIPROD	RZG	700,000	0	~ 0 %	LRZ	700,000	0	~ 0 %	exec assigned
90492	LGICTAMD	ICHEC	3,024,000	0	~ 0 %	ICHEC	3,024,000	0	~ 0 %	exec assigned
90494	MAESTRO	CINECA	2,000,001	0	~ 0 %	IDRIS	2,000,000	0	~ 0 %	exec assigned
90495	MIXTUDI	CINECA	3,750,000	0	~ 0 %	CINECA	552,000	0	~ 0 %	exec assigned
						FZJ	3,198,000	0	~ 0 %	exec assigned
90497	MUSIC	PDC	231,000	0	~ 0 %	IDRIS	231,000	0	~ 0 %	In setup
90498	NANOBIO-2	ICHEC	1,999,997	0	~ 0 %	HLRS	1,960,000	0	~ 0 %	exec assigned
						CINES	40,000	0	~ 0 %	exec assigned
90500	NR-NSNS-BHNS	FZJ	3,000,000	0	~ 0 %	CSC	3,000,000	0	~ 0 %	exec assigned
90501	NUWCLAY	CINES	1,500,000	0	~ 0 %	CINECA	1,500,000	0	~ 0 %	*not yet started*
90503	PETAHUB	CINECA	1,800,000	0	~ 0 %	PDC	853,200	0	~ 0 %	exec assigned
						NCSA	946,800	0	~ 0 %	exec assigned
90504	PHOTMAT	HLRS	5,062,500	0	~ 0 %	CINES	2,521,125	0	~ 0 %	exec assigned
						PSNC	2,541,375	0	~ 0 %	exec assigned
90505	PICKH	IDRIS	1,500,000	0	~ 0 %	CINECA	1,500,000	0	~ 0 %	exec assigned
90507	Planck-LFI	CSC	3,500,000	0	~ 0 %	CSC	3,500,000	0	~ 0 %	exec assigned
90509	RBflow-2	SARA	6,000,000	0	~ 0 %	RZG	3,450,000	0	~ 0 %	exec assigned
						SARA	2,550,000	0	~ 0 %	exec assigned
90512	SCW	CINECA	1,410,000	0	~ 0 %	LRZ	1,410,000	0	~ 0 %	exec assigned
90513	SIMONA	PSNC	599,994	0	~ 0 %	IDRIS	60,000	0	~ 0 %	exec assigned
						SARA	372,000	0	~ 0 %	exec assigned
						PSNC	168,000	0	~ 0 %	exec assigned
90514	SIVE-2	PDC	6,250,000	0	~ 0 %	EPCC	6,250,000	0	~ 0 %	In setup
90515	SMARC	RZG	700,000	0	~ 0 %	LRZ	700,000	0	~ 0 %	exec assigned
90517	SPIESM	PDC	3,750,000	0	~ 0 %	PDC	3,750,000	0	~ 0 %	In setup
90518	TanGrin	CSC	3,500,000	0	~ 0 %	EPCC	3,500,000	0	~ 0 %	exec assigned
90520	VIRonSAMs	HLRS	4,480,000	0	~ 0 %	CINECA	2,213,120	0	~ 0 %	exec assigned
						CINES	2,266,880	0	~ 0 %	exec assigned
90522	WESF	IDRIS	1,200,000	0	~ 0 %	CSC	252,240	0	~ 0 %	exec assigned
						RZG	947,760	0	~ 0 %	exec assigned
	Σ		90,620,013	0	~ 0 %		90,619,838	0	~ 0 %	

Figure 4: DPMDB view on DECI-7 projects

3.7 Application support and porting

3.7.1 *Objective*

Application support within DECI aims to facilitate a better understanding of the likely requirements of future users of the Tier-0 systems by collecting real use-case information about the needs and capabilities of scientific codes and about the differences between usage of national and European resources and facilities. By working on scalability and performance aspects of scientific codes, application support enables DECI to become a very effective instrument for providing a ramp from Tier-1 to Tier-0.

Moreover, the possibility of having applications experts from leading European HPC centres provide support to scientists that are interested in using the opportunity within DECI to experiment with new architectures is also highly valued.

In the course of 2011 two user questionnaires were sent to DECI researchers in order to investigate the impact of DECI infrastructure and get feedback on the quality of the provided services. The activities where DECI appeared to have had the greatest impact were in "Scaling codes to run on larger HPC systems" and in "Obtaining access to expert technical or application enabling support".

The statistics of these questionnaires was summarized in D4.3.1 [1] and suggested that the primary benefit which DECI delivers is to enable researchers to scale codes to run on larger HPC systems. The outcome of the questionnaires clearly emphasized that the applications enabling support which is provided with DECI is an important and highly regarded feature of the service.

Support for the first Tier-1 call under PRACE AISBL begins at a very early stage, when applicants are writing their proposal (advice on what machines to ask for, answering queries, etc.) and starts from the first day of opening the DECI call. Application experts from PRACE centres in close contact with researchers give them technical support in questions related to the proposal. The process of application support is a responsibility of PRACE-2IP WP7 task T7.2 and is described in more detail in WP7 deliverables.

3.7.2 *Structure*

Within PRACE-2IP application support is a responsibility of WP7 task T7.2. Application experts from each DECI site as well as other PRACE HPC centres that are not active in DECI but have expressed their wish to assist DECI applicants are identified for this task.

WP7 representatives participate in the monthly DECI video conferences so that all aspects of researcher support of particular projects can be discussed The specific WP7 DECI support task also holds teleconferences on monthly basis to coordinate the support activities of DECI projects between all PRACE HPC centres involved. More frequent teleconferences are being held during the first two months of the DECI allocation time period in order to ensure a smooth access and porting of the codes of DECI applicants to the PRACE systems.

T7.2 application support task assigns each DECI accepted project an expert from the home site (e.g. the site that gave support to the applicant for the DECI call) as well as an expert from the execution site (e.g. the site where the application is going to be executed). If the DECI application has multiple execution sites then it is assigned an expert from each execution site to help with porting of the code to the given system.

During the whole DECI process the work of T7.2 and WP2 is tightly coupled and dependent on each other. Figure 5: DECI process gives a clear indication how the whole process works.

The application expert from the home site has the following obligations:

1. Contact the PI (either by physically visiting him/her or by phone and finally by email if the PI is a returning DECI user). During the meeting the home site expert should explain the whole DECI procedure from the user point of view, inform him/her about final report obligations and specifics about the assigned execution site. Home site experts should also help the PI to apply for the national x.509 certificate.
2. Ensure that the execution site expert is informed about the PI and has necessary information to help the PI to port the code to the execution system.

It is important to note that home site expert is responsible to inform the PI about the services PRACE provides:

- Trouble ticket system and contact information
- Module environment
- PRACE common work environment – GPFS file system setup
- PRACE user documentation

The home site expert should also ensure that the DECI users have accepted the PRACE Acceptable User Policy before the accounts will be opened.

Execution site expert provides the PI with the PRACE login account to access the execution site and helps with porting of the code.

In addition to the home site and execution site experts, the T7.2 application support task also selects one expert known as the enabling expert for applications that have requested elaborate help with their codes in the DECI application form.

The enabling expert can be either the home expert or the execution expert who has a major and very important task to work with the PI in an intensive manner to help with

- performance
- scalability
- optimization
- parallelization

aspects of the code. The effort timeframe for this enabling task is estimated for each DECI application during the technical review of the proposal. The timeframe may vary between 1 to 6 months depending on the PI's specified request.

As soon as PI has access to the assigned execution system, the enabling expert contacts the PI and starts the work on the code. The progress of this work is updated in a dedicated workflow database and can be easily followed by all members of the T7.2 task.

Whether the PI requested enabling or not the execution of the project should start and the progress should be constantly monitored, e.g. by checking the CPU usage, and possible obstacles in the project should be worked out. If the project did not request help for enabling and intends to do only production runs then the progress of the project is monitored by the execution site expert. If the PI requested enabling then the project progress is monitored by the enabling expert as this expert is in close collaboration with the PI.

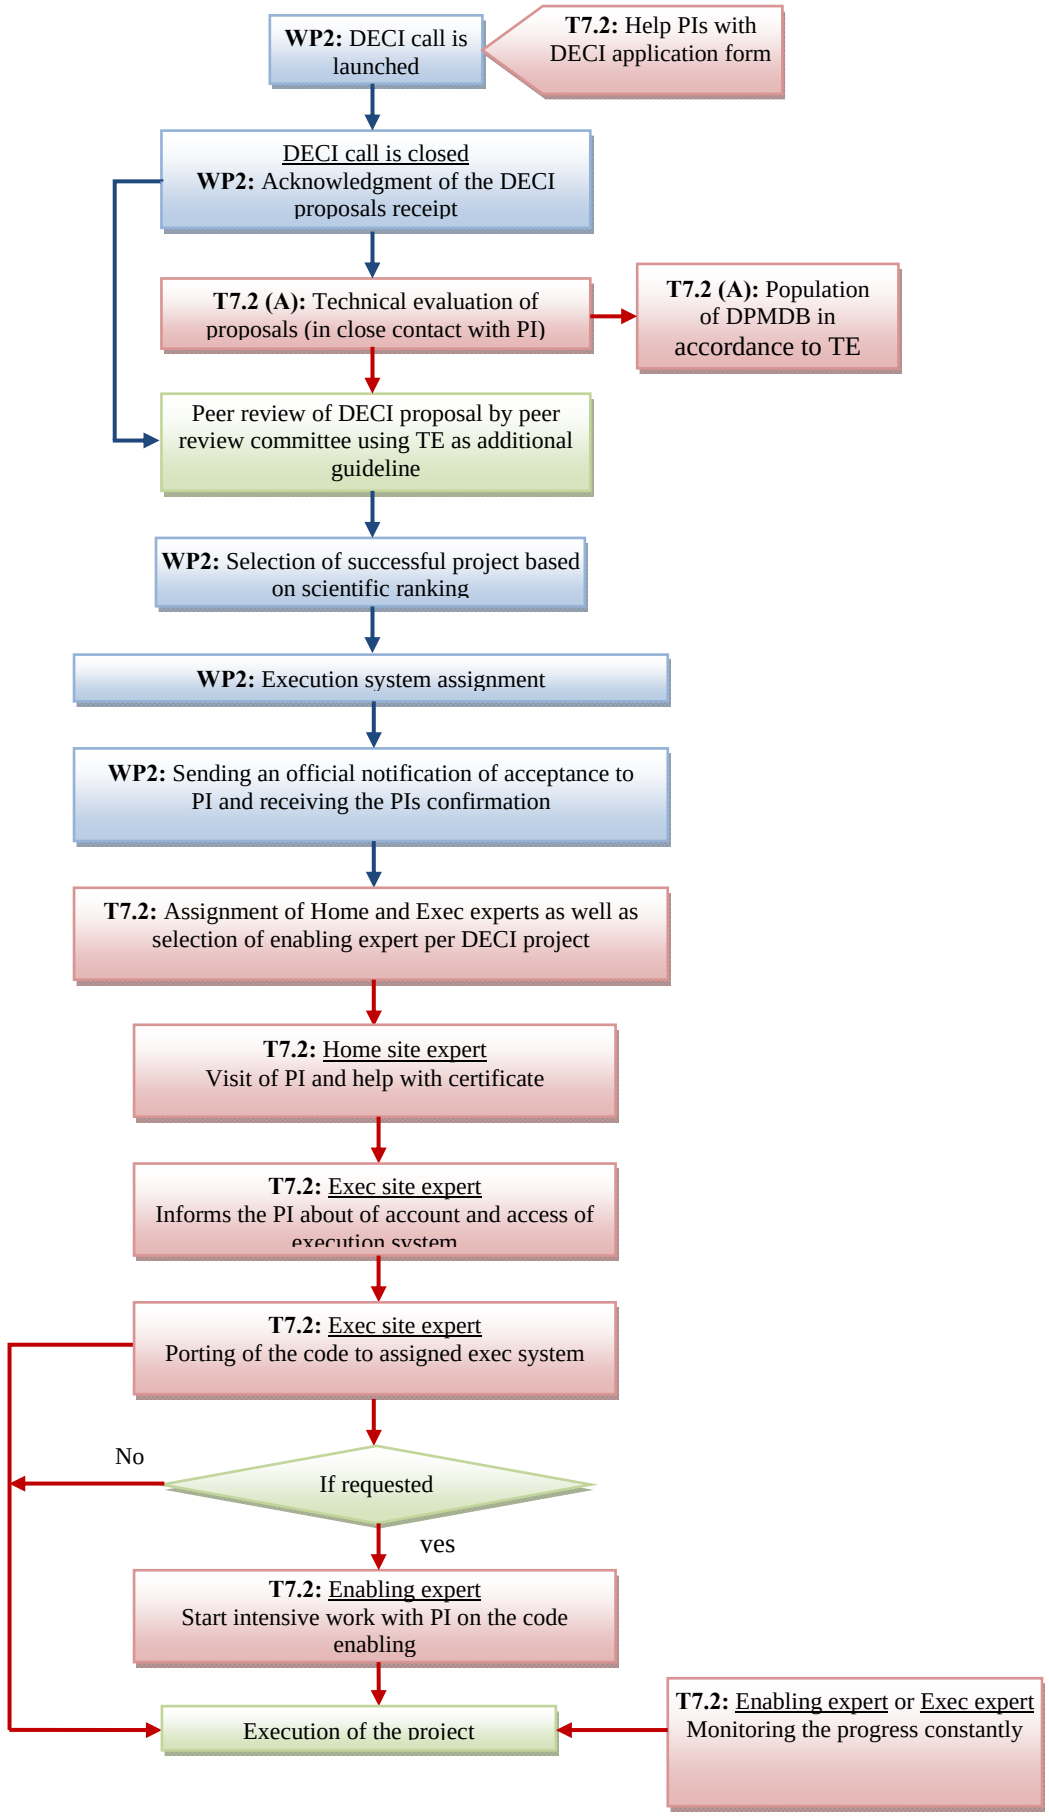

Figure 5: DECI process

3.8 Project follow-up

The progress of the DECI projects is followed by their home sites, and the status is recorded in DPMDB, see Figure 4: DPMDB view on DECI-7 projects. Status of enabling activity is recorded in the Best Practice Workflow (see the deliverable D7.2.1 and D7.2.2 of T7.2 for more details). Regular video conferences by WP2 are held every month to follow the course of DECI process and the progress of the DECI projects; this practice has previously been found to be very useful and effective in DEISA. The minutes of these meetings are stored in PRACE BSCW.

3.9 Project reports

DECI projects will submit a final scientific report when they have finished their project. These reports will be used for dissemination purposes within PRACE and for reviewing best practices, if necessary. They are collected by the home site representative of WP2, and stored in PRACE BSCW.

4 Comparison with the PRACE Tier-0 project access processes

4.1 Proposal submission

PRACE Tier-0 proposals are submitted electronically by means of an specially designed PRACE proposal submission system. This greatly simplifies the process both for the applicant and for the project staff. However, Tier-1 (DECI) proposals are submitted by email.

4.1.1 *Current DECI submission system*

Applicants apply to DECI by filling in an MS Word template document. As an example, the DECI-8 application form can be found at http://www.prace-ri.eu/IMG/doc/DECI8-PROPOSAL-ACRONYM.doc [2]. The completed form is then sent in to a dedicated email address and an acknowledgement is sent back manually. Once the form has been received, it is given a brief check and any immediate queries are referred back to the applicant.

Very soon after the closing date a definitive list of proposals is circulated amongst DECI staff and a designated centre is agreed for dealing with each proposal. The designated centre will complete the technical evaluation and is likely to become the home site for that project if this project is accepted. The application forms are then converted to PDF using a standard naming convention based on the given acronym and each form is filed within the BSCW in folders according to the designated home site. An entry is created in the DPMDB for each proposal and each home site then adds basic information (PI and Co-PI details, requested core-hours, etc.). All further administration is done via the BSCW until the point at which the projects are accepted when the DPMDB takes over again. However, the DPMDB contains links to the BSCW for each proposal, technical evaluation and scientific evaluation.

The MS Word template was originally based on the form used for Tier-0 applicants, although it has been modified for DECI use. Problems exist in using MS Word as different versions view the form differently and it is difficult to constrain applicants to filling in the form correctly. For example, many of the fields would be best presented using pull-down menus but these are difficult to implement within MS Word and so free text is allowed. Similarly several fields have “check boxes” to fill in, but there is no reliable way of requiring applicants to fill in the relevant boxes. Moreover, there is a large amount of information which needs to be transferred from the application forms to the DPMDB and technical evaluations, etc. This

presently has to happen by hand using a copy-and-paste approach or in some cases via manually re-typing the information. This can be both time-consuming and error-prone, especially in the case where many applications are received. Addresses can be particularly problematic to manipulate. To mitigate the problem of errors due to mis-typing, we implement a checking system where centres check each other's information. The technical and scientific evaluations similarly consist of MS Word forms which are filed in the BSCW.

4.1.2 *Alternatives for submission*

The above problems have already been solved via the system used within the HPC-Europa (TA) Transnational Access Visitor Programme. Here applicants apply via a web-based (PHP) system where personal information (address, phone number, gender, etc.) only has to be entered once for the lifetime of the whole HPC-Europa project and is then available within the on-line database to HPC-Europa staff. Application forms are completed on-line and can be completed in stages with details saved under a personal login. When the form is finally submitted and closed an email is sent to the applicant automatically with a reference number. Technical and scientific evaluations are then entered on-line via dedicated login-password combinations. All the information is then viewable on-line via various "cuts" of the data so that, for example, the projects can be listed against their technical and scientific evaluation scores or a spreadsheet of PI addresses can be generated automatically without staff having had to fill in any of these details themselves.

As well as much commonality between the HPC-Europa visitor programme and DECI, there are of course some differences between the two projects. In particular under HPC-Europa each application is associated with just one applicant (the "visitor"), whereas each DECI project is usually associated with multiple investigators. In addition within DECI we already have a database (the DPMDB) and so designing a suitable interface between the on-line application system and the DPMDB would be needed.

PRACE Tier-0 proposals are submitted using a similar system provided by CINES. This web-based submission system API is also written in PHP. The access is through the address https://prace-peer-review.cines.fr/. The access needs a registration and password. A research plan in PDF can be uploaded to the system. The research plan has a template in RTF. The technical and scientific evaluations are both done on-line as with HPC-Europa.

Adopting this technology for DECI has the potential to produce huge savings in terms of staff effort and would give the process and more professional feel than is possible via the present combination of email and MS Word.

4.2 Process for launching the calls

In order to improve the call process for future calls, a teleconference with AISBL, PMO, WP3, press and web team members was held on 10th November.

The following process for the call has been proposed as agreed in this meeting:

1. There should be a document which outlines what happens during a PRACE call.
2. A PRACE call should be driven by one person; the Driver
3. A PRACE call involved a team with different roles including: the Driver, the Approver, the resource representatives, the PMO, the Web and Press teams.
4. All roles should have a back up at all times.
5. Each PRACE call should be a plan, defined by the Driver.
6. All roles should be 'Informed' about the status with respect to the plan.

The call Approver is the PRACE Director. The Driver is a person from AISBL for the PRACE Tier-0 calls. The Driver for a DECI Tier-1 call is PMO. The Driver for a Joint Tier-0 & Tier-1 call is yet undefined; it will be decided later.

The PRACE AISBL will take responsibility for 'driving' the calls in the future but the next call, due in May 2012, is not yet defined and will need to be initiated in March 2012 and 'driven' some other way.

5 Conclusions

During the previous six years of its existence, the DECI programme and DECI processes have evolved in response to changing needs and the changing role which it plays in the HPC eco-system.

In analysing the incorporation of DECI into PRACE, and to make recommendations as to its future role in the HPC ecosystem, it is necessary to try to separate out a number of inter-connected strands and factors.

- Factors concerned with the expansion in the number of contributing partners from a fixed number of project partners (national HPC centres) participating in every call, to a much larger and more amorphous and fluid group of resource providers, participating on a call by call basis.
- Factors concerned with the much smaller average contribution made by any one country to the overall pool of resource. Under DEISA DECI, the average site hosted about five projects per call; the average for PRACE DECI to date is around two.
- Factors related with the heterogeneity of the HPC resources attached. Many more types of machines are included in DECI.
- Factors concerned with the much bigger spread of sizes of machines. The largest machines available via DECI are in the Top 20 system worldwide, the smallest machines are not even in the Top 500.
- Factors concerned with the desirability of closer collaboration in management and operation of calls between Tier-0 and Tier-1. This is largely a management decision concerning the role of Tier-1 in the HPC-ecosystem and whether its primary purpose is to provide a clear migration path and ramp from PRACE Tier-1 to PRACE Tier-0 resources or whether the primary role of DECI is to facilitate a resource sharing at the Tier-1 level between national providers.

At the present moment, it seems desirable to provide DECI with a more formal framework in which to operate. The process for becoming a Tier-0 partner and the rights and responsibilities and obligations that such as status confers are well-understood within the project. But the status, rights and responsibilities and obligations of Tier-1 partners are not nearly so clear. A PRACE optional programme would provide one possible way forward, and we recommend that this option be investigated by WP2.

6 Annex

The DECI-8 Proposal form and the Technical evaluation form are displayed in the next chapters.

6.1 DECI-8 Proposal form

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

APPLICATION FORM
DECI 8 call

General Information

Type of proposal	DECI-8 PROJECT
Start date	1st May 2012

Project Title:	
Project acronym:	
Research field	

Project leader (personal data and contact)

e-mail	
Gender	**Male/Female**
Title	
First name	
Last name	
Initials	
Date of birth	
Nationality	
Phone number	
Fax number	

Project leader (organisation and job title)

Job title	
Organisation name	
Group	
Department	
Address	
City	
Postal code	
Country	

Contact person for all correspondence (if different from above)

Name	
E-mail	

1

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

APPLICATION FORM
DECI 8 call

Organisation name	
Department	
Address	
City	
Postal code	
Country	
Job title	
Gender	
Title	
First name	
Last name	
Initials	
Date of birth	
Nationality	
e-mail	
Phone number	
Fax number	

Collaborators (add more as required)

Organisation name	
Department	
Address	
City	
Postal code	
Country	
Job title	
Gender	
Title	
First name	
Last name	
Initials	
Date of birth	
Nationality	
e-mail	
Phone number	
Fax number	

2

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

APPLICATION FORM
DECI 8 call

Summary of the project: (Maximum 500 words): If the project is successful this will be published on the PRACE website unless you mark it as confidential below. Please make this summary understandable to a general audience. (500 words)

3

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

APPLICATION FORM
DECI 8 call

Recent bibliographic references that are relevant to the project:

4

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

APPLICATION FORM
DECI 8 call

Computer resources requested

Total number of core hours requested: _____

Please explain how this figure was calculated, i.e. what is the number of planned simulations, number of cores, what platform is assumed, including performance figures, etc.:

Please describe how you foresee employing the PRACE infrastructure, for instance, do you plan to run multiple simulations concurrently on a single platform, do you wish to chain jobs, or run multiple codes as a workflow using one or more platforms from potentially different architectures, etc. Further, please state if you wish to employ grid middleware (GridFTP, UNICORE, DESHL, etc.).

5

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

APPLICATION FORM
DECI 8 call

In this section, we are interested in each of the codes you intend to run. Please provide the following details for ***each*** of your codes, employing a copy of this page to describe each code in turn.

Name of code: _____

Please rank your preferred target PRACE architecture
(1=first choice, 2=second choice, xxx, x=not suitable)

Class of machine	Ranking
IBM BlueGene/P	
IBM Power 6	
Cray XT4/5/6, XE6	
Cluster (Xeon/Nehalem+SGI/PowerPC/Westmere+Bullx)	
Cluster (GPU)	

Please give the motivation for your preferences above:
_____
_____
_____

Expected number of cores (please give both a minimum to make the feasible and the maximum the simulation can exploit): _____

Amount of memory per core (please give both the minimum to make the project feasible and the maximum amount the code can exploit): _____

Wall clock time for a single run of this code (min/max): _____
If max>12hours, does your simulation have a checkpoint/restart feature: **y/n**

Number of simulations to be run (min/max): _____

What applications and/or libraries does this code require: _____

Is this code I/O intensive ? **Yes/No**
If the answer is Yes, please describe your strategy concerning I/O (for example usage of I/O libraries, MPI I/O, netcdf, HDF5 or other approaches):
_____
_____

Amount of hard disk space required by this code (min/max) _____

6

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

APPLICATION FORM
DECI 8 call

In this section, we are interested in the overall data requirements for your project as a whole. Please provide the following information.

What is the amount of data to be transferred to the target platforms before production runs can start:

_____

What is the amount of data to be transferred from the target platform after all production runs are finished:

_____

Is it possible to start transferring data before all production runs are finished? **y/n**

Where will the data eventually be stored? (HPC facility, local storage at university, institute):

_____

If any planned runs require extensive amounts of memory, hard disk space, etc., please provide details on how the project (in terms of CPU hours) may be split across multiple platforms.

_____
_____
_____

7

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

APPLICATION FORM
DECI 8 call

Describe what work has already been done to develop the codes

This should include the following: describing the main algorithms, how they have been implemented and parallelized, and their main performance bottlenecks and the solutions to the performance issues you have considered. Please provide the name and version of all codes to be used in the project. For each code that needs to be optimized, please provide the details listed below.

1. Name and version.
2. Webpage and other references.
3. Licensing model.
4. Contact information of the code developers.
5. Your relationship to the code (developer, collaborator to main developers, end user, etc.).

8

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

APPLICATION FORM
DECI 8 call

Describe what enabling work will need to be completed before production runs can begin. Please make it clear what work will need to be done by your own group and what you are requesting to be done by PRACE staff. Please also make it clear if any enabling work can be done in parallel with production runs:

9

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

APPLICATION FORM
DECI 8 call

Discuss the routes that you will use for dissemination of the project and for any appropriate knowledge transfer. This should include any resources that you will be using to support this. (Maximum 500 words) :

Confidentiality

Is any part of the project covered by confidentiality? Yes/No

If YES, please give the reasons for confidentiality:

10

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

APPLICATION FORM
DECI 8 call

Do you have any other support for this application e.g. from your national funding council, the EC or international collaborations? Please give details of this below :

Please use the remainder of this form to give a detailed description of the project and its relevance for PRACE.

Please provide 3-4 pages, including:
1. scientific objectives
2. scientific and technical innovation potential
3. current profile and performance of code(s), including scalability, requirements on interconnect, I/O, architecture, clarification of how requested core-hours was calculated, are jobs independent, chained and/or workflows, etc.
4. computational objectives
5. specific benefits expected from PRACE
6. summary

11

6.2 Technical Evaluation form

Technical review

Proposal 2011-**PUT-ACRONYM-HERE** :

Tier 1 technical review

1. General information

1.1. Proposal

Centre undertaking technical evaluation:	
Name and email address of person undertaking technical evaluation:	
Proposal acronym :	
Project title (if given) :	
Research field :	
Proposal type :	Tier 1

1.2. Project leader (PI)

Name:	
Job title:	
Organisation:	

Organisation country:	
E-mail:	

2. Requested resources

2.1. Requested computer systems

Architectures requested	**Machine requested?**	**Suitability**		
IBM BlueGene/P	☐	Highly suitable Suitable Not suitable	☐ ☐ ☐	
IBM Power 6	☐	Highly suitable Suitable Not suitable	☐ ☐ ☐	
Cray XT4/5/6, XE6	☐	Highly suitable Suitable Not suitable	☐ ☐ ☐	
Cluster (Xeon/Nehalem + GPU/SGI/PowerPC/Westmere + GPU/Bullx+GPU)	☐ with GPUs? ☐	Highly suitable Suitable Not suitable	☐ ☐ ☐	Suggest specific cluster(s) (if appropriate):
No specific architectures requested on application form	☐	Highly suitable Suitable Not suitable	☐ ☐ ☐	

Comments on above including recommended machine:				
Has the code already been ported to other architectures? Yes ☐ / No ☐ **If yes, which ones?**				
IBM BlueGene/P ☐	IBM Power 6 ☐	Cray XT4/5/6, XE6 ☐	Cluster ☐ with GPUs? ☐	Other architecture ☐ Details:

2.2. Computer resources requested

Total core-hours	
Machine on which these hours are based	
Minimum core-hours necessary (if given or can be surmised)	
Expected no. cores	
Job wall clock time	
Total storage for duration of project (Gbyte)	

Maximum amount of memory per run (Mbyte)	
Maximum amount of memory per core (MByte)	

2.3. Description of the I/O strategy

a) Is I/O expected to be a bottleneck?

b) Implementation: I/O libraries, MPI I/O, netcdf, HDF5 or other approaches:

Does the project require grid middleware (e.g. GridFTP, DESHL, workflows)?	Yes/No
Comments:	

3 Project overview

3.1 Project targets

Are the targets proposed feasible within the allocation time from a technical viewpoint?	Yes/No
Comments:	

3.2 Enabling

a) Does the code require porting?	Yes/No
Comment:	
b) Does the project require further testing and optimisation? (If the answer is **Yes**, please explain and give an estimate of the amount of work necessary.)	Yes/No

Comment:
c) How much enabling is required overall (estimated)? (small < 1 mnth FTE, medium 1-3 mnth FTE, large > 3 mnth FTE)? Comment:

3.3. Additional comments from computer centres

4 Overall assessment

Overall result	
Project feasibility ***(0:** not feasible; **1:**needs some porting/enabling; **2:** no significant porting/enabling required)*	