SEVENTH FRAMEWORK PROGRAMME
Research Infrastructures

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

CAPACITIES

PRACE-1IP

PRACE First Implementation Project

Grant Agreement Number: RI-261557

D7.4.1
Applications and user requirements for Tier-0 systems

Final

Version:	1.0
Author(s):	Mark Bull (EPCC), Xu Guo (EPCC), Ioannis Liabotis (GRNET)
Date:	22.02.2011

Project and Deliverable Information Sheet

PRACE Project	**Project Ref. №: RI-261557**	
	Project Title: PRACE First Implementation Project	
	Project Web Site: http://www.prace-project.eu	
	Deliverable ID: D7.4.1	
	Deliverable Nature: DOC_TYPE: Report	
	Deliverable Level: PU	**Contractual Date of Delivery:** 28 / Feb / 2011
		Actual Date of Delivery: 28 / Feb / 2011
	EC Project Officer: Bernhard Fabianek	

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

Document Control Sheet

Document	**Title: Applications and user requirements for Tier-0 systems**	
	ID: D7.4.1	
	Version: 1.0	**Status:** For Approval
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2003	
	File(s): D7.4.1.docx	
Authorship	**Written by:**	Mark Bull, EPCC
	Contributors:	Xu Guo, EPCC Ioannis Liabotis, GRNET
	Reviewed by:	Vincent Ribaillier, CNRS/IDRIS Thomas Eickermann, FZJ
	Approved by:	Technical Board

Document Status Sheet

Version	Date	Status	Comments
0.1	01/02/2011	Draft	
0.2	04/02/2011	Revised draft	
0.3	08/02/2011	Draft for internal review	
0.9	18/02/2011	Revised after review	
1.0	22/02/2011	For Approval	

Document Keywords

Keywords:	PRACE, HPC, Research Infrastructure, Requirements, Systems, Applications, Users, Surveys

Table of Contents

List of Figures

List of Tables

References and Applicable Documents

[1] PRACE Project: http://www.prace-project.eu
[2] HP-SEE Project Deliverable D4.1, "Target Applications Analysis".
[3] PRACE Preparatory Phase Deliverable D6.1, “Identification and Categorisation of Applications and Initial Benchmarks Suite”.
[4] PRACE Preparatory Phase Deliverable D6.2.2, “Final Report on Application Requirements”.
[5] TOP500 Supercomputing Sites, http://www.top500.org

List of Acronyms and Abbreviations

BSC	Barcelona Supercomputing Center (Spain)
CEA	Commissariat à l’Energie Atomique (represented in PRACE by GENCI, France)
CINECA	Consorzio Interuniversitario, the largest Italian computing centre (Italy).
CSC	Finnish IT Centre for Science (Finland)
DEISA	Distributed European Infrastructure for Supercomputing Applications. EU project by leading national HPC centres.

EPCC	Edinburgh Parallel Computing Centre (represented in PRACE by EPSRC, United Kingdom)
FZJ	Forschungszentrum Jülich (Germany)
GB	Giga (= 2^{30} ~ 10^9) Bytes (= 8 bits), also GByte
GB/s	Giga (= 10^9) Bytes (= 8 bits) per second, also GByte/s
GFlop/s	Giga (= 10^9) Floating point operations (usually in 64-bit, i.e. DP) per second, also GF/s
GHz	Giga (= 10^9) Hertz, frequency =10^9 periods or clock cycles per second
GPGPU	General Purpose GPU
GPU	Graphic Processing Unit
HLRS	High Performance Computing Center Stuttgart (Germany)
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing
HPF	High Performance Fortran
HP-SEE	High-Performance Computing Infrastructure for South East Europe's Research Communities, Project Acronym
ICHEC	Irish Centre for High-End Computing (Ireland)
IDRIS	Institut du Développement et des Ressources en Informatique Scientifique (represented in PRACE by GENCI, France)
IPB	Institute of Physics, Belgrade (Serbia)
KTH	Kungliga Tekniska Högskolan (represented in PRACE by SNIC, Sweden)
LEF	LINPACK Equivalent Flop/s
LINPACK	Software library for Linear Algebra
LRZ	Leibniz Supercomputing Centre (Garching, Germany)
MB	Mega (= 2^{20} ~ 10^6) Bytes (= 8 bits), also MByte
MB/s	Mega (= 10^6) Bytes (= 8 bits) per second, also MByte/s
MPI	Message Passing Interface
NTNU	Norwegian University of Science and Technology (Trondheim, Norway)
OpenMP	Open Multi-Processing
PFlop/s	Peta (= 10^{15}) Floating-point operations (usually in 64-bit, i.e. DP) per second, also PF/s
PGAS	Partitioned Global Address Space
PLEF	Peta-LEF
PRACE	Partnership for Advanced Computing in Europe; Project Acronym
PRACE-PP	PRACE Preparatory Phase Project
PSNC	Poznan Supercomputing and Networking Centre (Poland)
SEE	South Eastern Europe
SNIC	Swedish National Infrastructure for Computing (Sweden)
TFlop/s	Tera (= 10^{12}) Floating-point operations (usually in 64-bit, i.e. DP) per second, also TF/s
Tier-0	Denotes the apex of a conceptual pyramid of HPC systems. In this context the Supercomputing Research Infrastructure would host the Tier-0 systems; national or topical HPC centres would constitute Tier-1
TLEF	Tera-LEF
UiB	University of Bergen (Norway)
UPC	Unified Parallel C
VRC	Virtual Research Community

Executive Summary

This is the initial deliverable for Task 7.4 (Applications Requirements for Tier-0 Systems). It reports the results of surveys carried out of current PRACE partners' HPC systems, the applications running on them, and of current/potential users of the PRACE infrastructure. The questions asked in the systems and applications survey are largely the same as were asked in a previous PRACE survey, carried out in 2008, and so changes in findings can be observed. The results of an applications survey from the South Eastern European HPC community are included as an Annex to this deliverable.

The principal findings from the PRACE surveys were:

- 28 systems are represented in the systems survey, representing just over 3.0 PFlop/s of peak computing power.
- Compared to the 2008 PRACE-PP survey, the compute power available across the PRACE partners has increased by a factor of 3.7, with almost all this increase being a result of increased numbers of cores, rather than increased power per core.
- Around a quarter of CPU cycles on these systems go unused.
- As the number of cores in systems has increased, so too has the number of cores used by jobs on these systems, but at a slower rate, so that the fraction of the machine used by the average job has decreased since 2008.
- PRACE partners are supporting about 50% more users per system than in 2008.
- While the popular scientific and I/O libraries, debuggers and performance analysis tools are present on many systems, it is clear that they are not universally available.
- For each system, we requested an application survey return for all applications which consumed more than 5% of the CPU cycles on that system. This resulted in 93 survey returns, representing 57 different applications.
- The most widely used applications are largely the same ones as in the 2008 survey.
- Compared to the 2008 survey, there has been an increase in the proportion of applications using C or C++, such that the balance is now approximately half C/C++ and half FORTRAN.
- MPI remains by far the most popular parallelisation technique, though there has been a modest increase in the number of codes using mixed-mode MPI/OpenMP.
- There were no applications reported as using any of the PGAS family of APIs.
- 411 responses to the user survey were collected.
- Over 50% of users have their own application codes, and consider themselves as developers rather than end users.
- The most commonly used shared codes are mainly from the areas of Computational Chemistry and Molecular Dynamics.
- There is a general desire from users to increase the scalability of their applications, but their ambitions are relatively modest.

- Over one-third of users do not fully understand the scalability issues of their applications and about 15% expressed a desire for assistance from PRACE to solve this problem.
- Over one-third of users require more than 2GB of memory per core for their application: this is likely to become a significant problem on future systems.
- Almost 50% of users would be prepared to use a different application if it meant more scalability.
- Requirements for disk space and for the amount of data to be transferred off the system vary widely, over more than three orders of magnitude.
- Users are slightly more likely to use C or C++ than FORTRAN as their main development language
- MPI and OpenMP are by far the most widely used parallelisation methods.
- Use of Grid middleware and workflow systems is rather low.
- Just over 50% of users had considered applying for PRACE resources, but of those that had not, over half were unaware of the possibility.
- A large majority of users thought that the potential future diversity of architectures would make applying for PRACE resources more attractive, and that smaller, stepping-stone resources (i.e. Tier-1) would be a helpful route to Tier-0 access.

1 Introduction

This document contains the results of surveys which were carried out as part of the work of Task 7.4 (Applications Requirements for Tier-0 Systems) of the PRACE First Implementation Phase Project (PRACE-1IP) [1]. The surveys were:

- a survey of HPC systems operated by PRACE partners,
- a survey of the most important applications running on these systems, and
- a survey of people who are currently, or may potentially be, users of the PRACE infrastructure.

These surveys fulfill several purposes:

- to understand which applications may in future be used on PRACE Tier-0 systems,
- to assess the requirements of these applications for CPU cycles, data storage, libraries, tools, etc.,
- to understand how these applications are making use of current systems operated by PRACE partners, and
- to understand the requirements of users of the PRACE Tier-0 infrastructure.

Section 2 of this document contains a summary of the results of the systems and applications surveys. These surveys were aimed at providing a snapshot of current usage of HPC systems operated by PRACE partners. The questions asked are largely the same as were asked in the previous PRACE survey, carried out in 2008, and reported in [3]. The earlier survey is referred to in the remainder of this document as the 2008 PRACE-PP survey. Because we collected essentially the same data in the two surveys, separated by a gap of just over two-and-a-half years, we are able to observe some changes over the intervening period.

Section 3 contains a summary of the results of the user survey. This survey was aimed at providing a comprehensive picture of usage patterns, behaviour and requirements of current and potential users of PRACE Tier-0 systems. In terms of the questions asked, there is some overlap with the survey conducted by WP6 in the PRACE-PP project, and reported in [4], but this survey is more comprehensive. A number of PRACE-1IP workpackages and tasks contributed questions to the current survey, including WP4 (HPC Ecosystem Relations), WP6 (Technical Operation and Evolution of the Distributed Infrastructure) and Task 7.1 (Applications Enabling for Capability Science). This document contains a summary of the results of the survey, but it is expected that the data will be used by the contributing tasks/workpackages for more detailed analysis. The survey sample was self-selecting: each partner was asked to contact any users who might potentially have an interest in using PRACE systems, and invite them to take part.

Section 4 contains a summary of the results of the surveys, and draws some conclusions, including some recommendations on tools and libraries.

Annex A of this deliverable contains the results of an application survey from the South Eastern European HPC community. At around the same time of conducting the PRACE-1IP WP7 applications survey, the HP-SEE (http://www.hp-see.eu) project conducted an individual applications survey with a project defined questionnaire. The HP-SEE applications survey was targeted to the applications that are proposed within the project and was conducted in the period of October to November 2010. Since the results of this survey may also be of interest to readers, this deliverable presents the main findings. More details on the HP-SEE project, as well as on the scope of the questionnaires, are presented in Annex A. Note that scientific users in PRACE member countries that participate in HP-SEE, were also given access to the PRACE-1IP WP7 applications survey and may have responded to the WP7 user survey as well. The results provided in Annex A are a summary of a public deliverable of the HP-SEE project [2].

2 System and Application Surveys

2.1 Overview

For the systems and applications surveys, to a significant extent we have asked many of the same questions as were asked in the 2008 PRACE-PP survey. This allows us to observe any changes in the intervening period. We collected 28 system survey returns, which represent the major systems of PRACE partners and other large national systems.

Each partner was then asked to complete an applications survey for each application on their system that accounted for more than 5% of the utilisation. We collected 92 application survey returns, representing 57 distinct applications.

As in the 2008 PRACE-PP survey, we use the notion of LINPACK Equivalent Flop/s (or LEF) to measure the consumed machine time: for example an application which uses 10% of the cycles on a machine with a 100 Tflop/s performance rating for the LINPACK benchmark would have a rating of 10 TLEF.

The responses were collected between 23rd November 2010 and 17th January 2011. The survey period for each system was at least three months, and the end of this period was as some point during 2010. For each system, data were collected for the top applications running on that system over at least a three month period. This information was collected using two online survey forms; one for the system and another for applications running on that system.

The system survey then asked the following:

- Generic details of the system: Name, manufacturer, model, processor type, clock rate, memory, configuration of system (cores per chip, chips per node, etc), I/O configuration, cache, interconnect system.
- Performance figures: R_{max}, R_{peak}, availability, utilisation.
- The use of the system: job sizes, scientific areas, number of users.
- System software: scientific libraries, compilers, performance analysis tools, I/O libraries, parallel debugging tools.
- The top applications: applications using > 5% of the available cycles. For each of these an application survey was expected.

The applications survey asked for the following information:
- Generic information on the application: name, description, authors.
- Scientific areas covered.
- Languages and libraries: languages used, parallelisation techniques, lines of code, libraries required.
- Usage: utilisation percentage of the application on the system in question, job size distribution.

The surveys were implemented using SurveyMonkey (www.surveymonkey.com) and the results were obtained in the form of Excel spreadsheets. These spreadsheets will be made available for use by other Workpackages in PRACE-1IP as required.

2.2 System survey results

Table 1 lists the systems for which a survey response was received, the HPC centre and country responsible for their operation, and the manufacturer and model of each system. The systems represent 23 centres from 12 countries.

Name	*Centre*	*Country*	*Manufacturer*	*Model*
Akka	SNIC	Sweden	IBM	Blade Centre HS21
Babel	IDRIS	France	IBM	Blue Gene/P
Baku	HLRS	Germany	NEC	SX-9
Bw.Grid Cluster	HLRS	Germany	IBM	Blade Center HS21 XM
Halo2	PSNC	Poland	Sun Microsystems	Sun Blade 6048 Constellation
HECToR Phase2a	EPCC	UK	CRAY	XT4
HECToR Phase2b	EPCC	UK	CRAY	XT6
Hexagon	UiB	Norway	Cray	XT4
HLRB-II	LRZ	Germany	SGI	SGI Altix 4700
Huygens	SARA	Netherlands	IBM	p575 cluster
Hydra	PSNC	Poland	HP	HP Blade System, HP c7000
Jugene	FZJ	Germany	IBM	Blue Gene/P
Juropa	FZJ	Germany	Sun-ParTec Cluster	Sun Blade 6048 system
Laki	HLRS	Germany	NEC	NEC HPC-144 Rb-1
Lindgren	KTH	Sweden	Cray	XT6m
Louhi	CSC	Finland	Cray	XT4/XT5
MareNostrum	BSC	Spain	IBM	PowerPC
Nautilus	PSNC	Poland	IBM	IBM QS22
Njord	NTNU	Norway	IBM	p575
Nova	PSNC (WCSS)	Poland	ACTION (Actina)	ACTION (Actina Solar)
PARADOX	IPB	Serbia	Intel	Xeon cluster
SP6	CINECA	Italy	IBM	IBM pSeries P575 Cluster
Stallo	Univ. of Tromsø	Norway	HP	BL 460c
Stokes	ICHEC	Ireland	SGI	Altix ICE 8200EX
Supernova	PSNC (WCSS)	Poland	Hewlett-Packard	Cluster Platform 3000 BL 2x220
Titan	Univ. of Oslo	Norway	Sun Microsystems	x2200
Titane	CEA	France	BULL	BULL Novascale R422 (CPU/GPU)
Vargas	IDRIS	France	IBM	pSeries 575

Table 1: List of systems for which a survey response was returned

Table 2 shows the peak (R_{peak}), and LINPACK (R_{max}) performance (as reported in the Top500 list) of each of the systems, together with the total number of cores in the system. The total power of systems is 3.07 Pflop/s peak, and 2.52 Pflop/s achieved LINPACK, from 573,360 cores. Compared to the 2008 PRACE-PP survey, the total computational power (as measured by R_{max}) available from all the surveyed systems has increased by a factor of about 3.7. The average R_{max} per core, however, has increased by less than 10%.

Name	*R_peak (Tflop/s)*	*R_max (Tflop/s)*	*Cores*
Jugene	1026.8	845.3	294912
HECToR Phase2b	366.7	274.7	44544
Juropa	212.0	187.9	17664
Babel	139.3	119.3	40960
HECToR Phase2a	113.1	95.1	12288
Titane	97.5	86.1	54816
SP6	101.1	78.0	5376
Louhi	102.0	76.5	10832
Lindgren	92.5	75.1	11040
MareNostrum	94.2	63.8	10240
Stallo	31.9	59.9	5632
HLRB-II	62.3	56.5	9728
Vargas	67.4	52.8	3584
Laki	62.7	50.8	5600
Huygens	62.6	48.9	3328
Akka	53.8	46.0	5376
Halo2	59.0	43.0	6912
Hexagon	51.1	40.6	5552
Stokes	41.0	36.6	3840
Supernova	40.9	34.7	3840
Titan	40.0	32.0[1]	4800
Bw.Grid Cluster	38.8	29.3	3424
Njord	22.5	19.0	3000
Nautilus	30.0	18.0	2016
Baku	19.7	17.3	192
Nova	19.0	16.0	2016
Hydra	12.2	8.6	1176
PARADOX	6.3	5.3	672
Total	**3066.0**	**2517.2**	**573360**

Table 2: Peak and Linpack performance of surveyed systems.

Figure 1 shows the percentage availability and the percentage utilisation of the available cycles for the 21 systems which reported this data. Availability for these systems was between 86% and 100%, with a mean of 97%. Utilisation was between 32% and 92% with a mean of 77%, which is higher than the mean of 71% reported in the 2008 PRACE-PP survey. The low figure of 32% for HECToR Phase2b can be attributed to the fact that the system had only been recently installed. Of the 2.28 LINPACK equivalent Pflop/s (PLEF) represented by these systems, 1.65 PLEF, or 72.4%, were actually consumed by applications.

[1] Estimated value based on 80% of R_{peak}

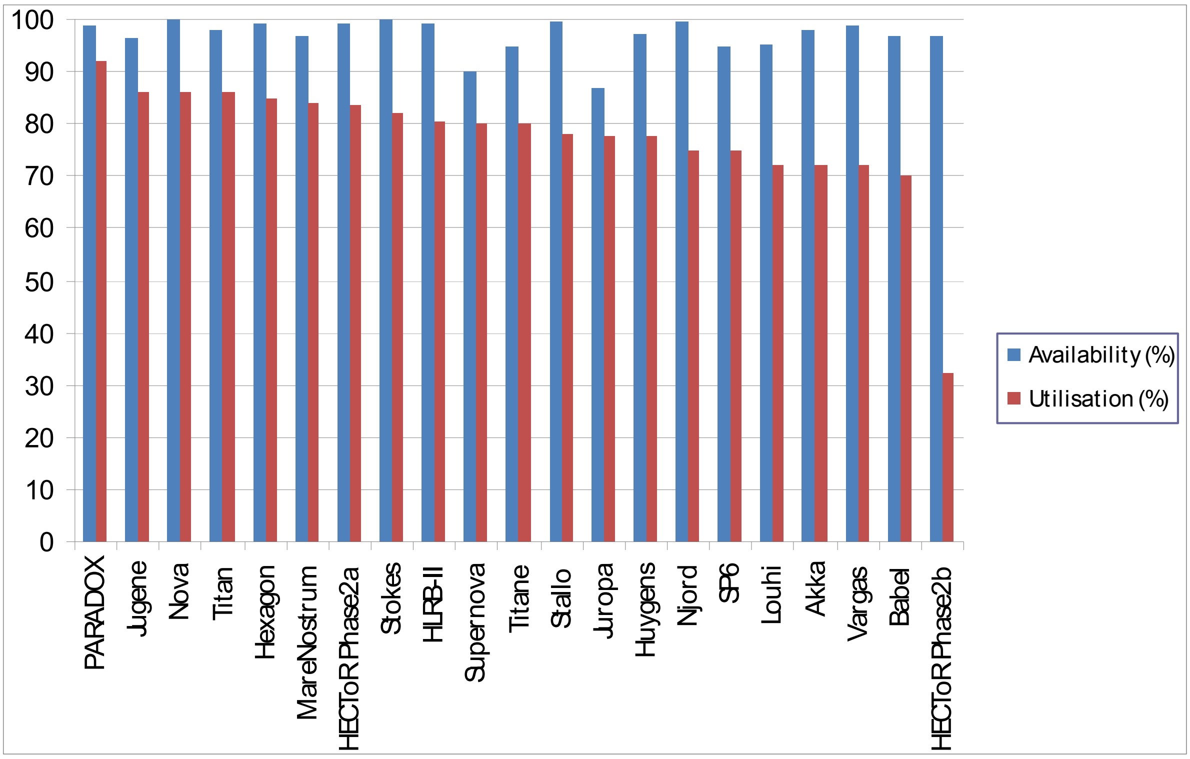

Figure 1: System availability and utilisation.

Figure 2 shows the job size distribution of the utilised cycles on each system for five ranges of job size: up to 128 cores, 128-512 cores, 513-2048 cores, 2049-8192 cores and more than 8193 cores. Note that the distribution is expressed as a percentage of the utilised cycles, not as a percentage of submitted jobs. A wide range of behaviour is observed: some systems run only large jobs, some only small jobs and some a more even spread across the ranges. Compared to the 2008 PRACE-PP survey, we observe a significant shift towards jobs using larger numbers of cores.

To further understand how the systems are utilised, we computed a mean job size for each system, by assuming that all the jobs in each range are on average the midpoint of the range (and that jobs in the >8193 range are assumed to be of size 16348, which may not be very accurate). We then divided this mean job size by the number of cores in the system, to obtain a metric which approximately represents the fraction of the system occupied by the average job. This metric is shown in Figure 3. The fraction of the system occupied by the average job varies from 1%, to 15%. This shows that the way machines are used varies widely: some systems are divided very finely between lots of small jobs, whereas others mostly run a small number of large jobs. Compared to the 2008 PRACE-PP survey, these figures are generally lower, which indicates that even though job sizes in terms of number of cores have increased, this increase has not kept pace with the number of cores in the machines, so that machines are now typically utilised by a larger number of concurrent jobs.

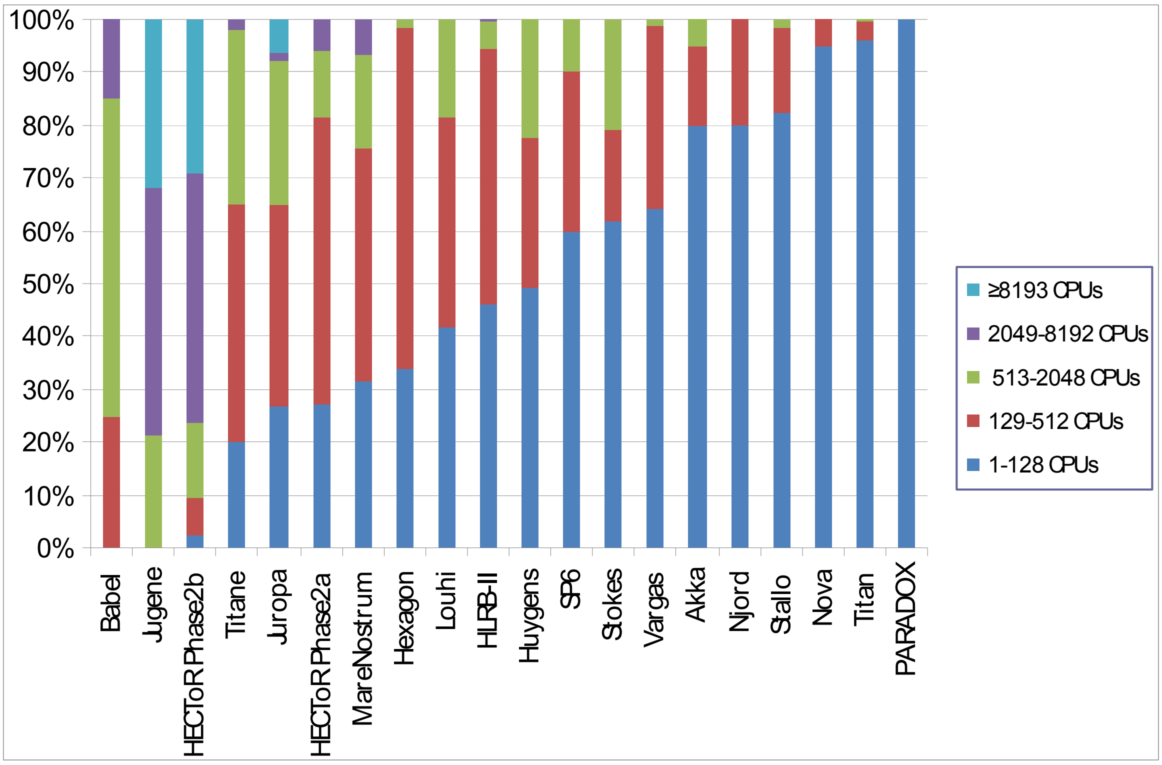

Figure 2: Job size distribution by system.

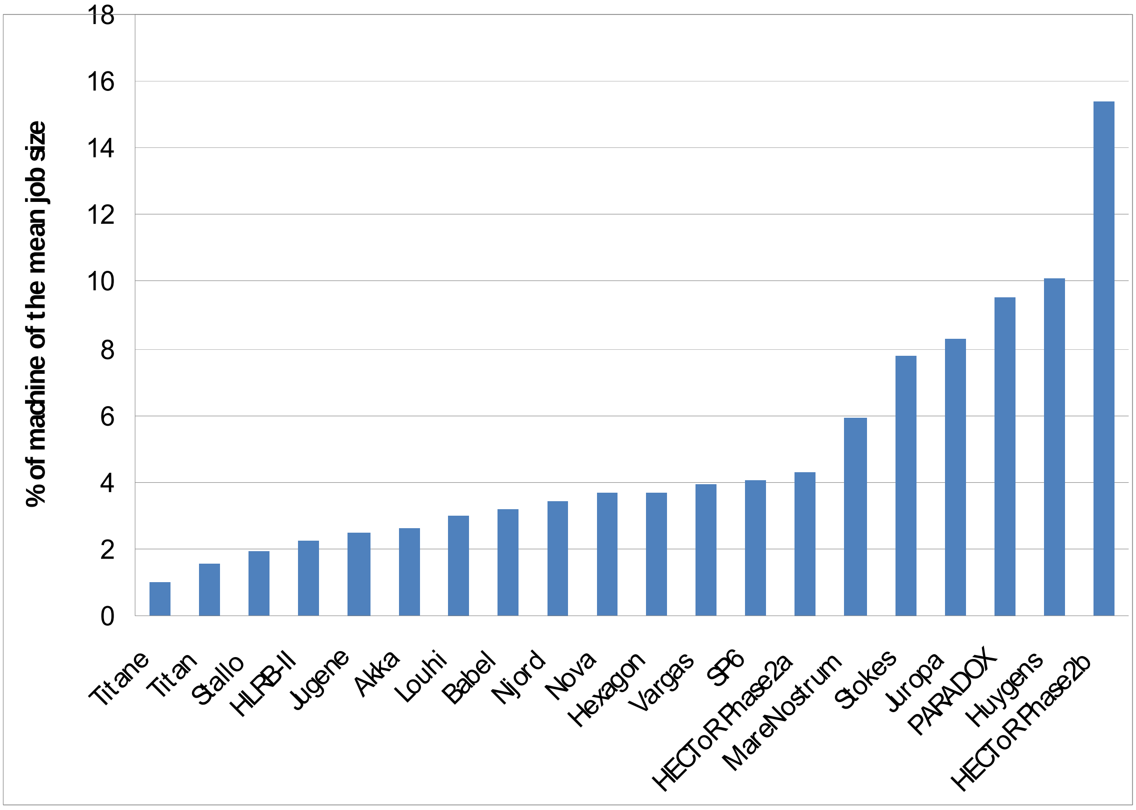

Figure 3: Mean job size as a percentage of system size.

Figure 4 shows the job size distribution aggregated across all systems. Compared to the 2008 PRACE-PP survey, the fraction of LEFs consumed by jobs with fewer than 129 cores has decreased from 37% to 22%, while the fraction consumed by jobs of over 2048 cores has increased from 26% to 40%.

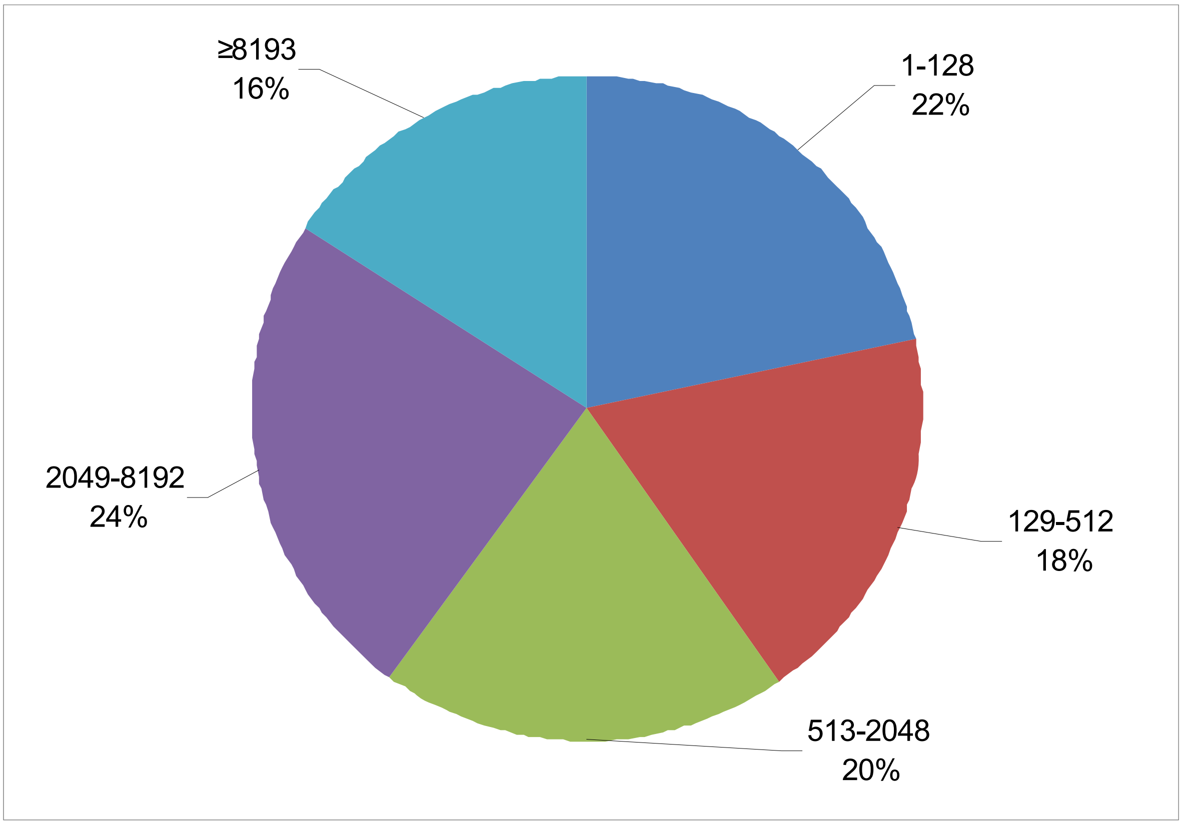

Figure 4: Aggregated distribution of LEFs by job size.

Figure 5 shows the distribution of cycles used by scientific area for 19 of the 28 systems. Only a few systems are dedicated to a small number of scientific areas: most systems have substantial usage from a number of different scientific areas.

Figure 6 shows the aggregated distribution of LEFs used in the different scientific areas across all the systems. Particle Physics accounts for over one quarter of the LEFs. The next largest areas are Computational Fluid Dynamics, Condensed Matter Physics and Computational Chemistry. Compared with the 2008 PRACE-PP survey, the proportion of cycles consumed by Particle Physics and Computational Fluid Dynamics is higher, while the proportion consumed by Condensed Matter Physics and Computational Chemistry is lower. It must be remembered, though, that the current survey figures represent a different set of systems/centres than those of the 2008 survey.

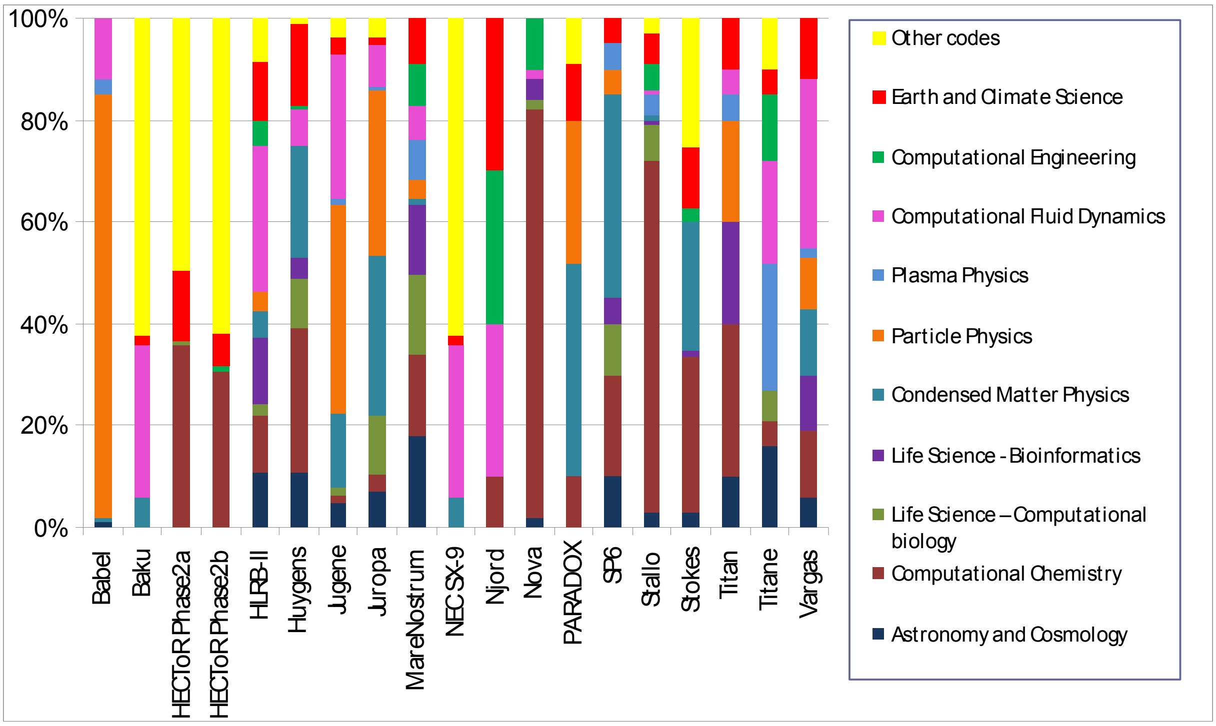

Figure 5: Scientific area distribution by system.

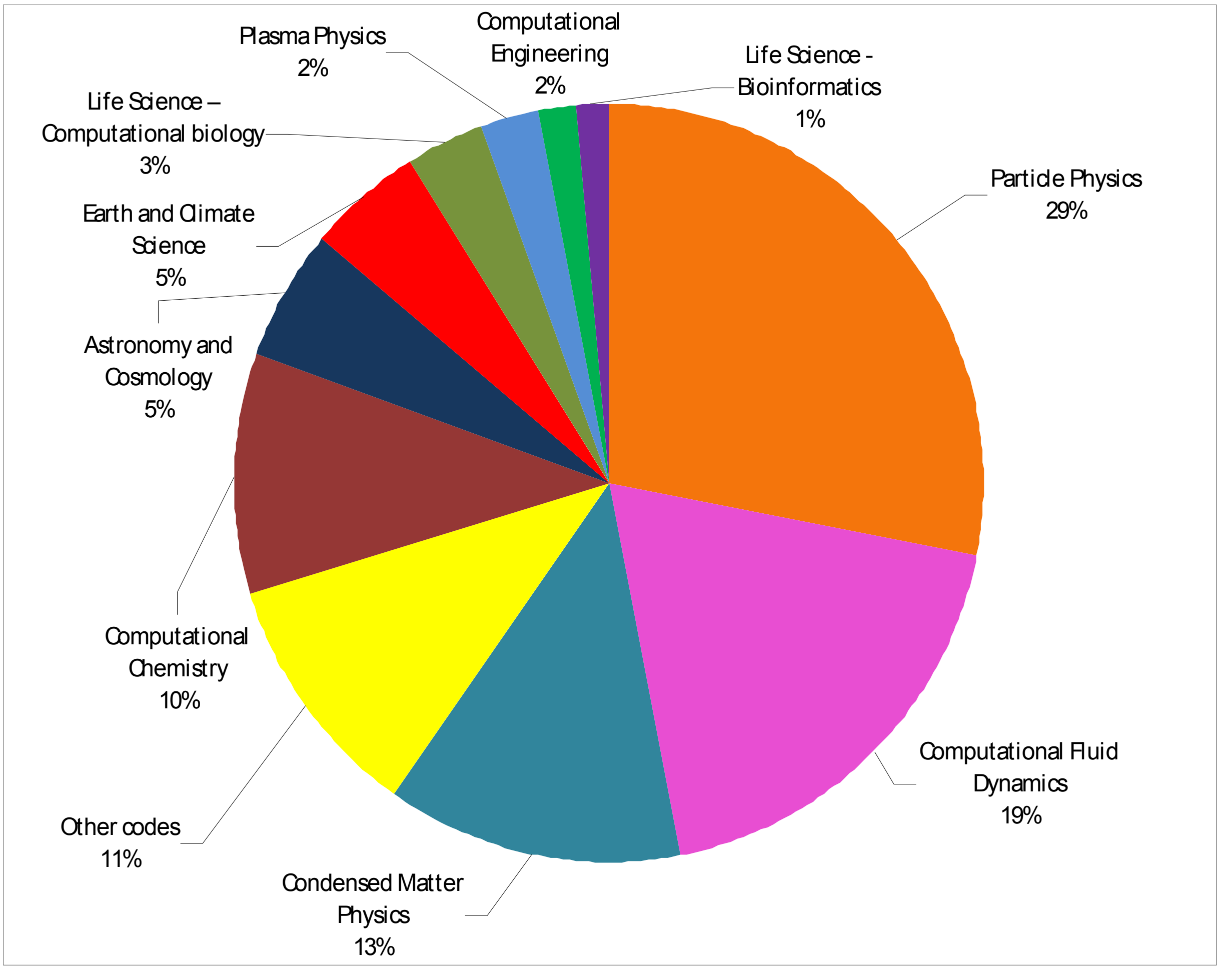

Figure 6: Aggregated distribution of LEFs by scientific area.

Twenty-two systems were able to report the number of users: this is shown in Figure 7 together with the available compute power in GLEF per user (obtained by dividing the system

R_{max} by the number of users). The number of users per system varies widely, from 30 to 2000. The total number of users on these 22 systems was 7875, giving an average of 358 users per system, an increase of 50% from the 2008 PRACE-PP survey. The compute power per user also shows a very wide variation, from 35 GLEF (PARADOX) to almost 2.5 TLEF (Jugene). On average (over all users) each user has access to 284 GLEF of compute power. This represents an increase of a factor of about 2.5 since the 2008 PRACE-PP survey. (In practise, the compute power is not evenly shared between users: observations suggest that in many cases a small number of users are responsible for using a high percentage of cycles on a given system.)

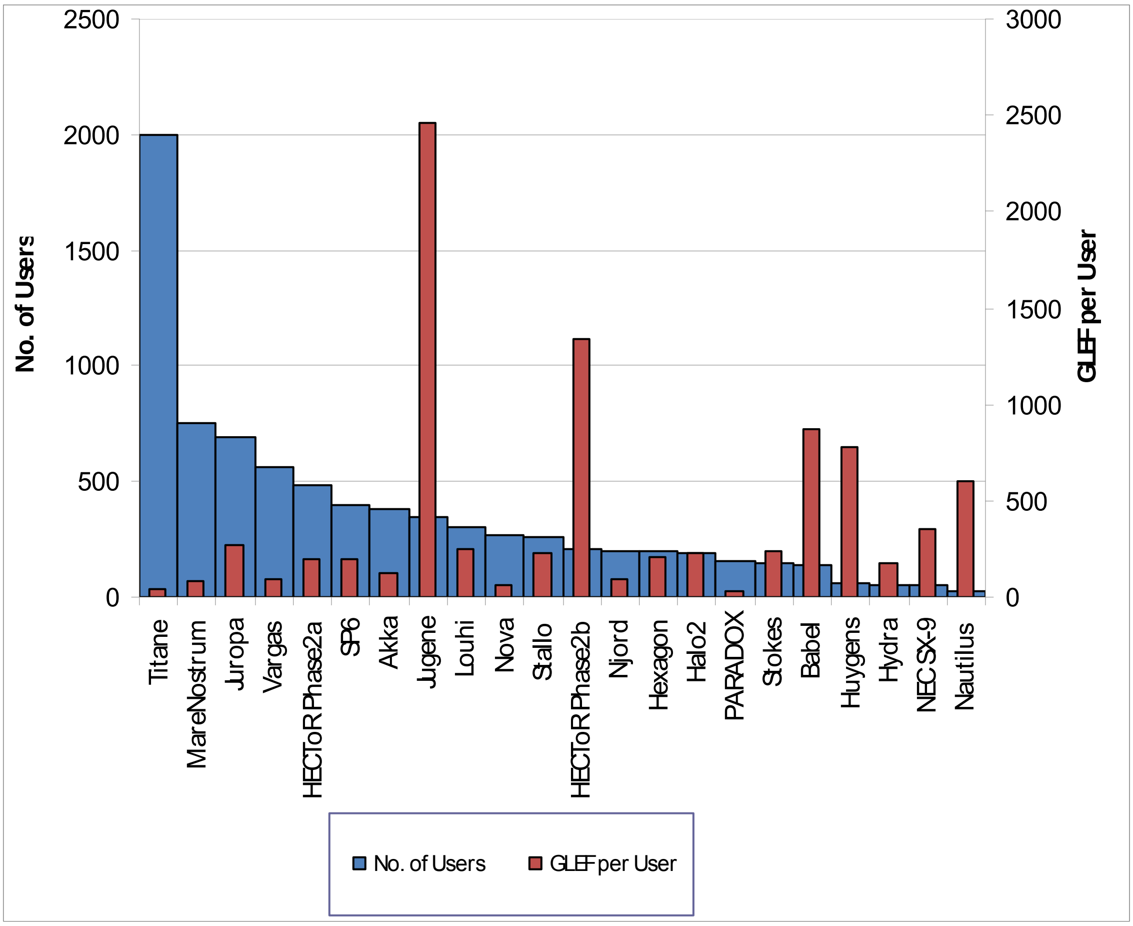

Figure 7: Number of users and GLEF per user.

Figure 8 shows the distribution of compilers installed on the systems. Gnu compilers are the most widely available, followed by Intel and PGI. Figure 9 shows the distribution of scientific libraries across the systems. FFTW and dense linear algebra libraries (LAPACK and versions of BLAS) are quite widely available, other libraries are present on fewer than 50% of the systems. Of course, it is often possible for users to install libraries themselves where needed, so a central installation in a system is not always required.

Figure 10 shows the distribution of I/O libraries: HDF5 and NETCDF are widely available. The installed MPI libraries are shown in Figure 11: there is a wide spread of different versions, which is not surprising as this library is quite closely tied to the interconnect hardware. Since the interface is so well standardized, there is little impact of this on the user.

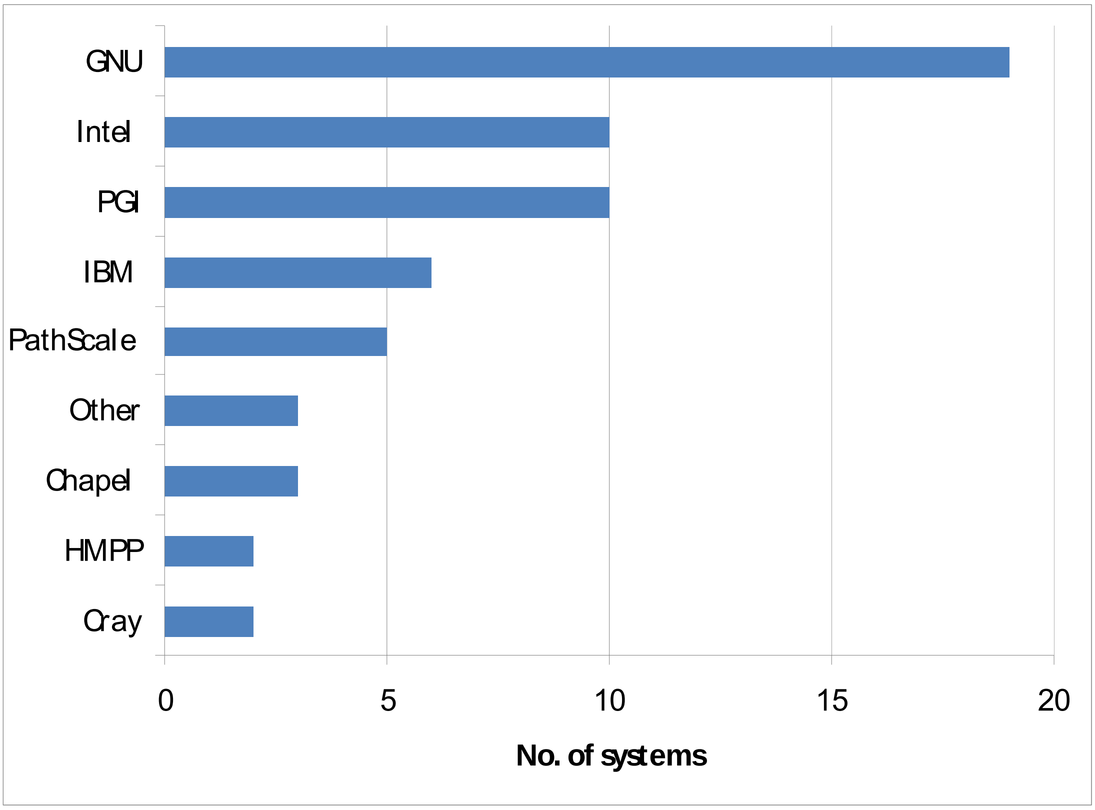

Figure 8: Installed compilers.

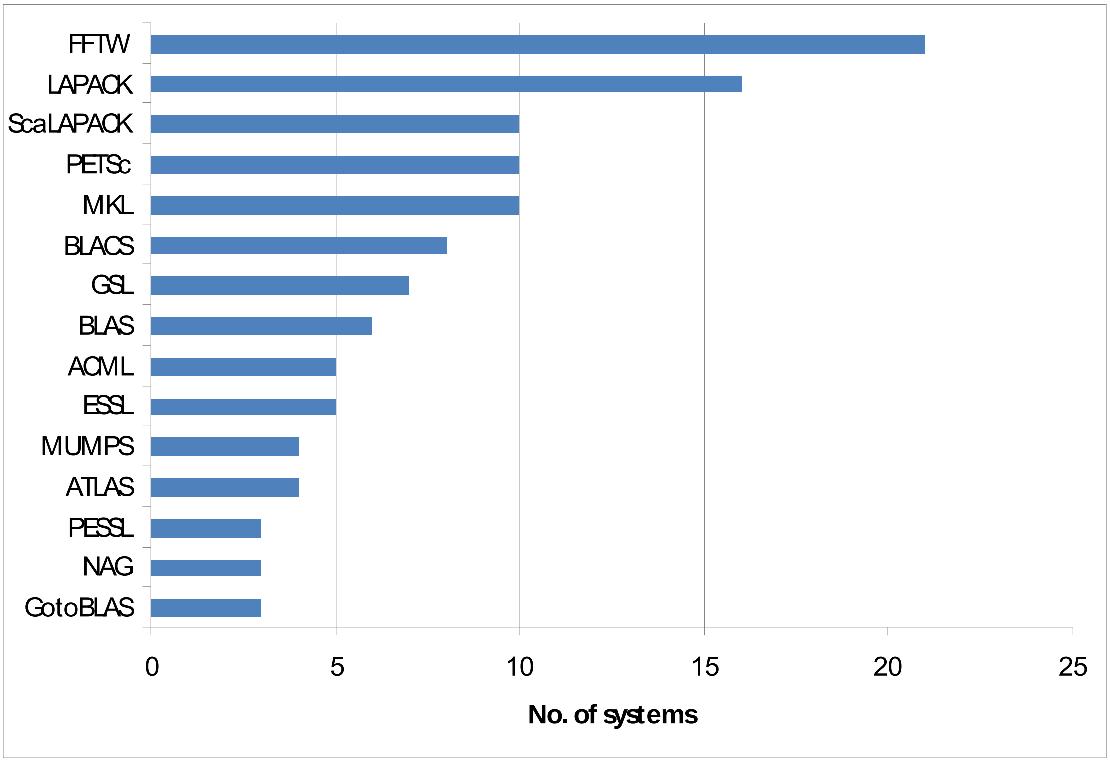

Figure 9: Installed scientific libraries.

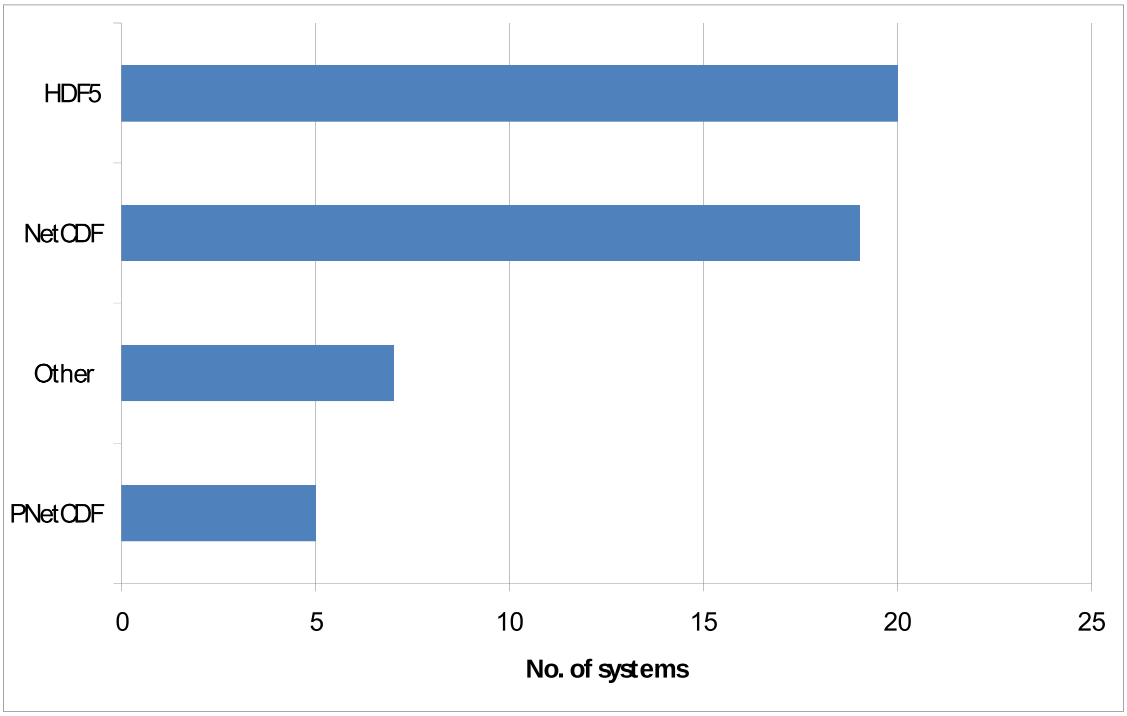

Figure 10: Installed I/O libraries.

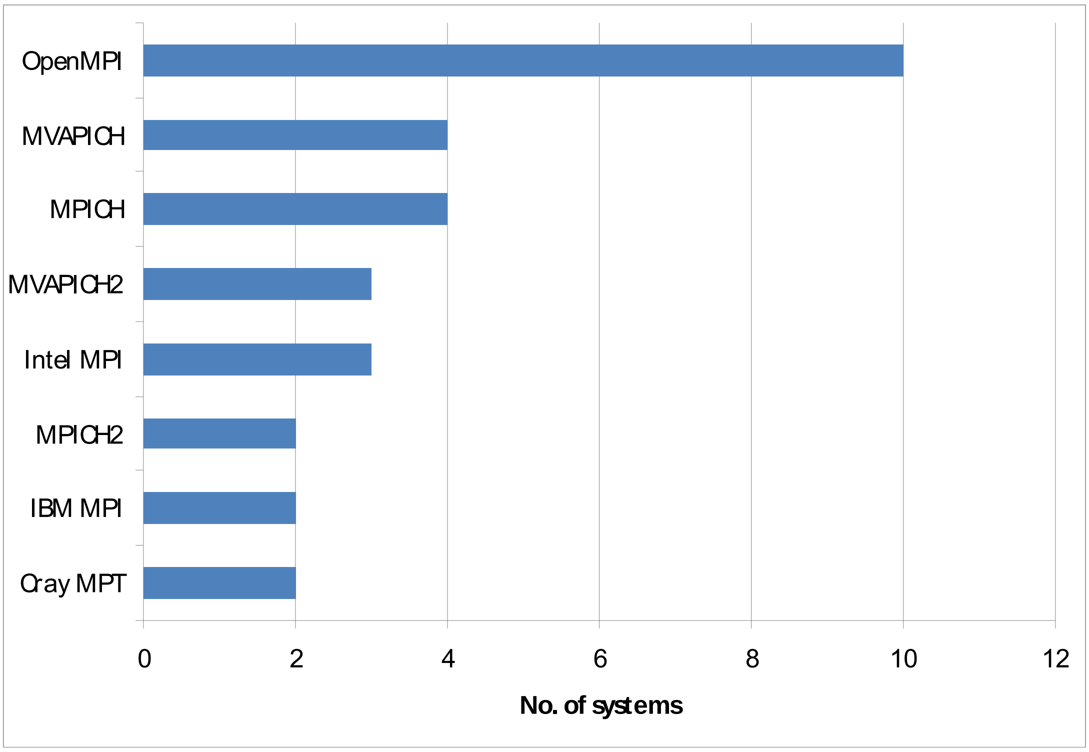

Figure 11: Installed MPI libraries.

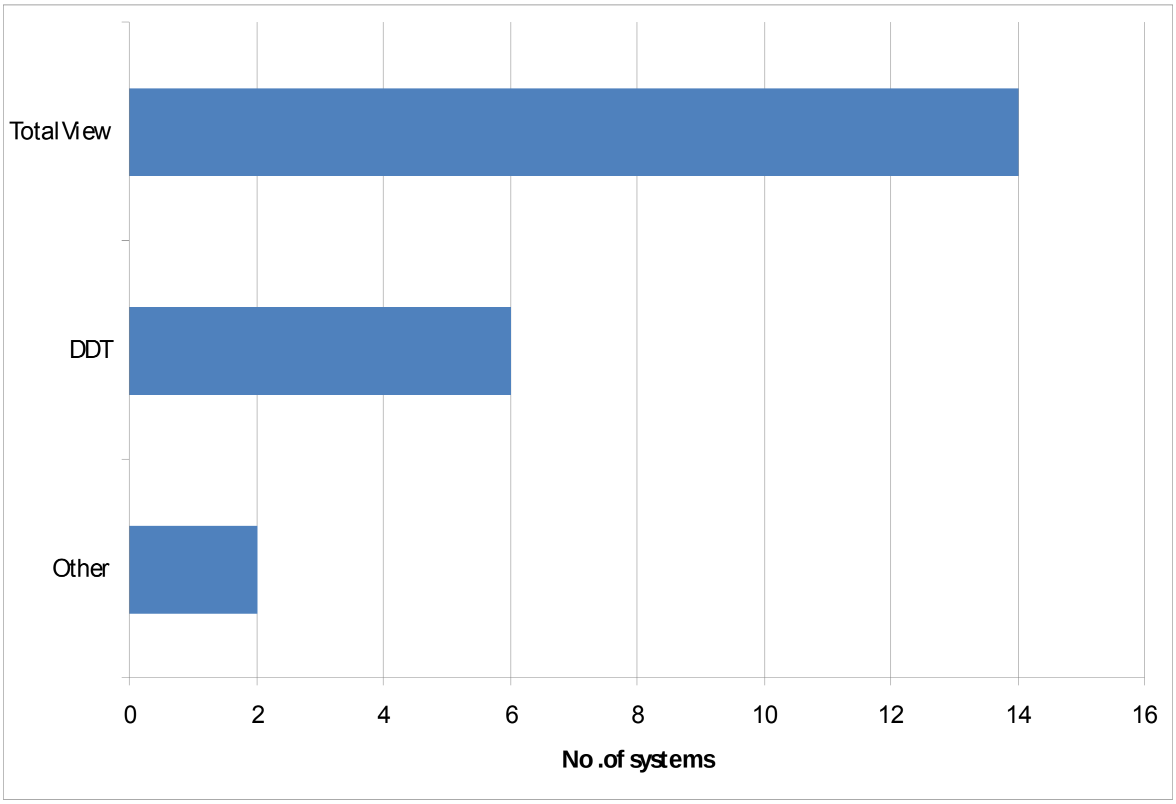

Figure 12: Installed debuggers.

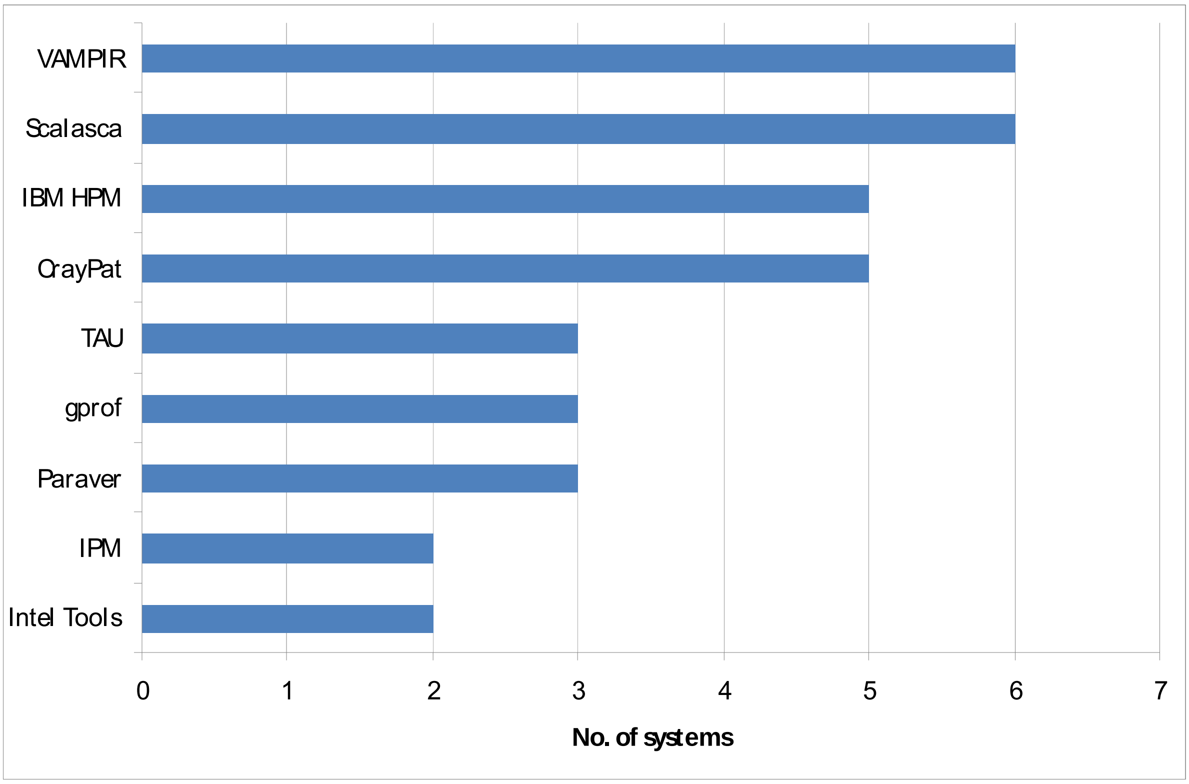

Figure 13: Installed performance analysis tools.

Figure 12 shows the distribution of installed debuggers: Totalview is by far the most popular, though not universal, and a small number of systems have DDT. Figure 13 shows the installed performance tools. There results show a lot of diversity: no tool is reported as present on more than 25% of the systems.

2.3 Application survey results

A total of 93 applications surveys were returned across the 28 systems: these come from 57 distinct applications. Table 3 lists these applications, ranked by the number of cycles consumed (as GLEFs, obtained as the product of the fraction of the machine time used with the R_{max} rating of the system) and the number of systems from which the application was reported. A significant number of applications survey returns did not include the utilisation figure: for the calculations in this table we have assumed 5% usage in these cases, as this was the minimum usage criterion for inclusion in the survey. The total usage of these applications in Table 3 is 786 Linpack-equivalent Tflop/s (31%) out of a total of 2517 Linpack-equivalent Tflop/s usage reported from the 28 systems.

Of the 57 applications, 42 were reported as being heavily (i.e., at least 5%) used on one system only, 7 on two systems and only seven (VASP, GROMACS, NAMD, Quantum Espresso, Gaussian, Gadget2 and CPMD) on three or more systems. The most widely used codes were VASP and GROMACS, which were reported as being used on eight different systems each. Compared with the 2008 PRACE-PP survey, Quantum Espresso, Gaussian and Gadget2 have joined this list of popular codes, while DALTON and AVBP have dropped out.

Application	*Total reported LEFs (Tflop/s)*	*No. of systems*
Thermodynamics with 2+1+1 dynamical flavors	137.8	1
Quantum Monte Carlo Studies of Strongly Correlated Systems	102.0	1
Lattice Boltzmann simulations of emulsification	81.1	1
QCD Simulations with Light, Strange and Charm Quark Flavors	63.4	1
CP2K	63.1	2
VASP	56.9	8
Gromacs	39.5	8
2 Flavour QCD with non-perturbative precision	30.2	1
NAMD	27.8	7
Quantum Espresso	19.7	3
Gaussian	14.0	3
Gadget2	11.9	3
Unified Model	10.8	2
GPAW	9.6	1
RAMSES	8.6	2
Bifrost	8.6	1
Enzo 1.5	7.8	1
CPMD	7.0	3
Wuppertal clover RHMC	6.0	1
SU3_Ahiggs	4.6	1
Papero	4.4	1
mglet	4.0	1
Dalton	3.8	1
ASH (Anelastic Spherical Harmonic)	3.7	1

Application	*Total reported LEFs (Tflop/s)*	*No. of systems*
tmLQCD	3.7	2
AMBER	3.6	2
POP	3.3	1
WRF	3.2	1
ADF	3.0	1
SES3D	2.8	1
ABAQUS	2.8	2
SSIM	2.7	1
AVBP	2.6	1
MESONH	2.6	1
NEMO	2.6	1
Fluent	2.6	2
User Code 4	2.5	1
User Code 6	2.5	1
User Code 8	2.5	1
User Code 9	2.5	1
COAMPS	2.2	1
OpenFOAM	1.7	1
SPEEDUP	1.7	1
GAMESS	1.3	1
ATHENA	1.1	1
Fnpeaks	0.9	1
User Code 1	0.9	1
User Code 2	0.9	1
User Code 3	0.9	1
octopus	0.7	1
Schroedinger	0.4	1
ETA	0.3	1
Continuous Time Quantum Monte Carlo	0.3	1
eqrfim	0.3	1
WRF-NMM	0.3	1
Total	**785.8**	

Table 3: Applications usage.

Application name	*System*	*Cores*	*1-128*	*129-512*	*513-2048*	*2049-8192*	*≥8193*
2 Flavour QCD	Juropa	17664		26	74		
ABAQUS	Nova	2016	100				
ABAQUS	Stokes	3840	95	5			
Amber	Hydra	1176	100				
ASH	Vargas	3584	2	98			
ATHENA	PARADOX	672	100				
AVBP	Vargas	3584	13	87			
Bifrost	Hexagon	5552	5	94	1		
COAMPS	Halo2	6912		100			
CP2K	HECToR	44544		4	7	9	80

Application name	*System*	*Cores*	*1-128*	*129-512*	*513-2048*	*2049-8192*	*≥8193*
	Phase2b						
CP2K	Huygens	3328	91	9	1		
CPMD	Nova	2016	100				
CPMD	Vargas	3584	100				
CTQMC	PARADOX	672	100				
Enzo 1.5	SP6	5376	10	70	20		
eqrfim	PARADOX	672	100				
ETA	PARADOX	672	100				
Fluent	Halo2	6912	100				
Fluent	Hydra	1176	100				
Fnpeaks	Nautilus	2016	100				
Gadget2	SP6	5376	10	80	10		
Gadget2	Nautilus	2016	100				
GAMESS	Nova	2016	100				
Gaussian	Halo2	6912	100				
Gaussian	Nova	2016	100				
GPAW	Louhi	10832	33	51	16		
GROMACS	HLRB-II	9728	79	21			
GROMACS	Huygens	3328	100				
GROMACS	Louhi	10832	68	31	2		
GROMACS	SP6	5376	50	45	5		
LAMMPS	Njord	3000	44	56			
LAMMPS	Njord	3000	44	56			
Lattice Boltzmann	Jugene	294912			1	62	37
MESONH	Vargas	3584	3	97			
mglet	Njord	3000	92	8			
NAMD	Louhi	10832	6	76	18		
NAMD	PARADOX	672	100				
NAMD	Vargas	3584	5	95			
NEMO	Vargas	3584	71	29			
OpenFOAM	Njord	3000	99	1			
Papero	Huygens	3328	82	18			
POP	Huygens	3328	27	9	63		
QCD Simulations	Jugene	294912			23	50	27
Quantum ESPRESSO	Jugene	294912	5	15	80		
Quantum Espresso	SP6	5376	50	45	5		
Quantum Monte Carlo	Juropa	17664	8	38	54		
RAMSES	Babel	40960			8		92
RAMSES	Vargas	3584	39	61			
Schroedinger	Hydra	1176	100				
SES3D	HLRB-II	9728	95	5			

Application name	System	Cores	1-128	129-512	513-2048	2049-8192	≥8193
SPEEDUP	PARADOX	672	100				
SSIM	Njord	3000	100				
SU3_Ahiggs	Louhi	10832	15	28	57		
Thermodynamics	Jugene	294912			27	44	29
tmLQCD	Babel	40960			69	30	1
tmLQCD	Vargas	3584	46	54			
Unified Model	Halo2	6912		100			
Unified Model	HECToR Phase2a	12288	59	41			
VASP	HECToR Phase2a	12288	48	37	14	2	
VASP	Stokes	3840	85	15			
VASP	Titan	4800	100				
VASP	Louhi	10832	59	38	3		
VASP	Huygens	3328	87	13			
WRF-NMM	PARADOX	672	100				
Wuppertal clover RHMC	Babel	40960			2	51	47

Table 4: Job size distribution as a percentage of jobs

Table 4 shows the job size distribution as a percentage of jobs in 5 ranges of core counts for those applications where this data was reported (Note that some application names have been truncated here). Many applications run on relatively modest core counts, while only a few exploit large numbers of cores.

Figure 14 shows the distribution of applications by scientific area. Note that some applications are used in more than one area: in this chart they are counted more than once.

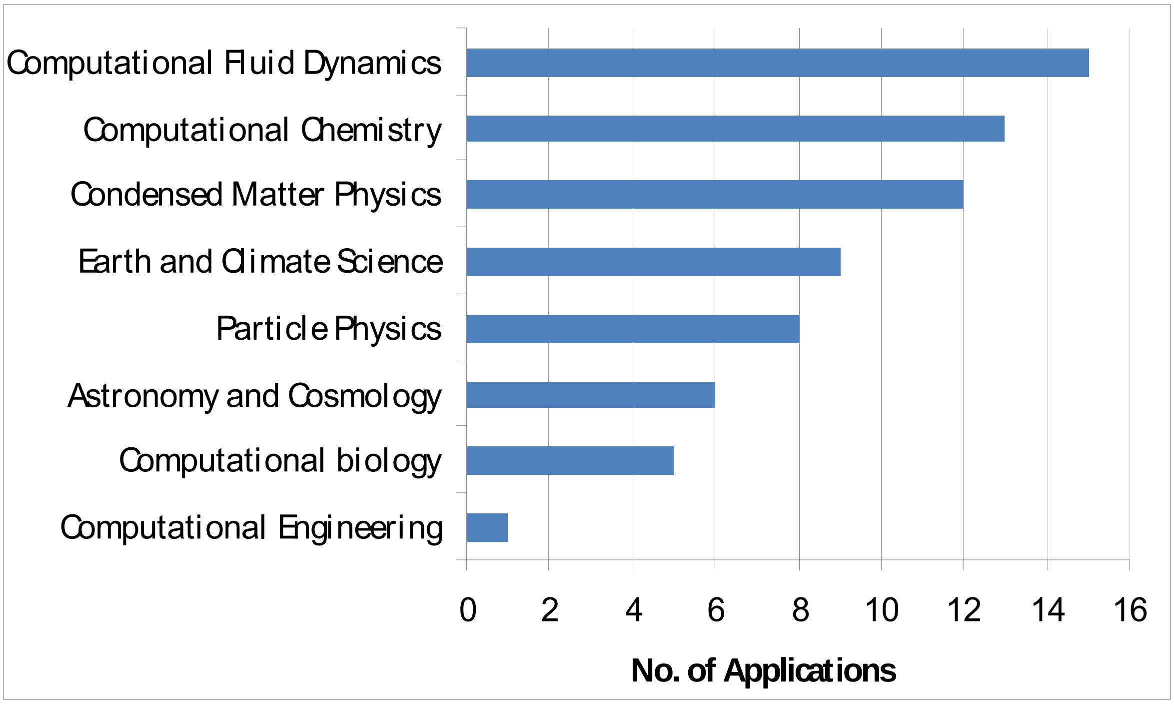

Figure 14: Distribution of applications by scientific area

Figure 15 shows the base languages utilised by the 57 distinct applications (note that some applications use more than one language). Compared with the 2008 PRACE-PP survey, we see an increase in the proportion of applications using C or C++ compared to those using Fortran: in this survey the balance was close to 50:50.

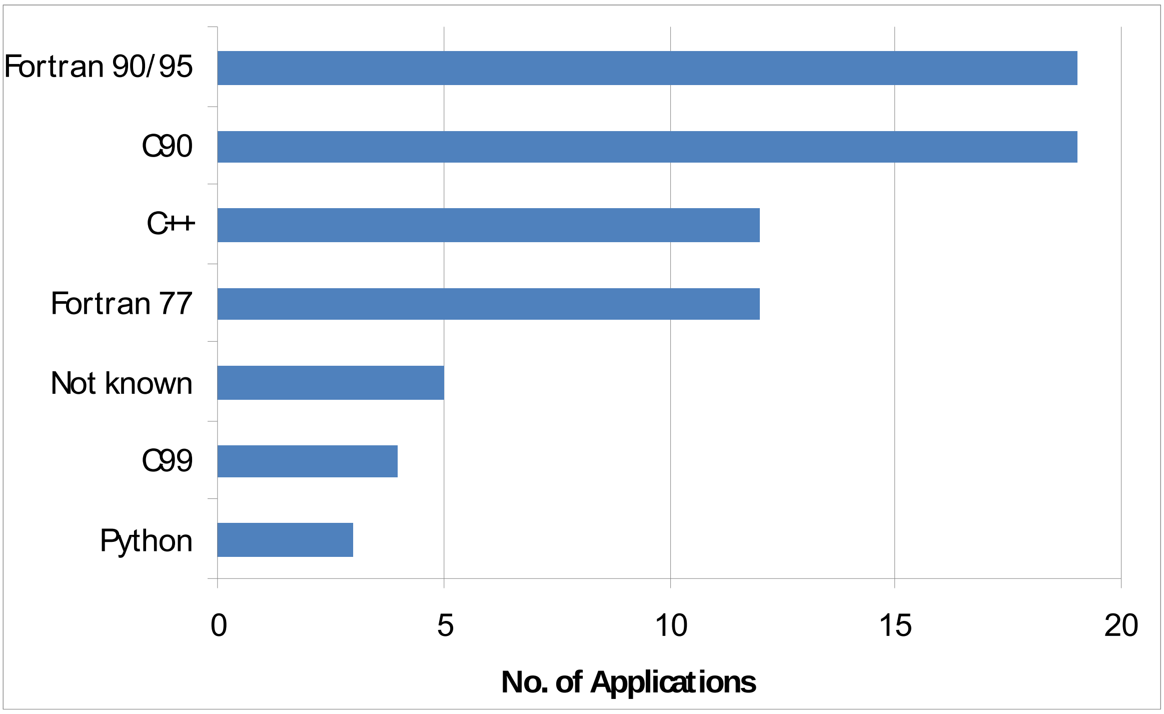

Figure 15: Base language utilisation.

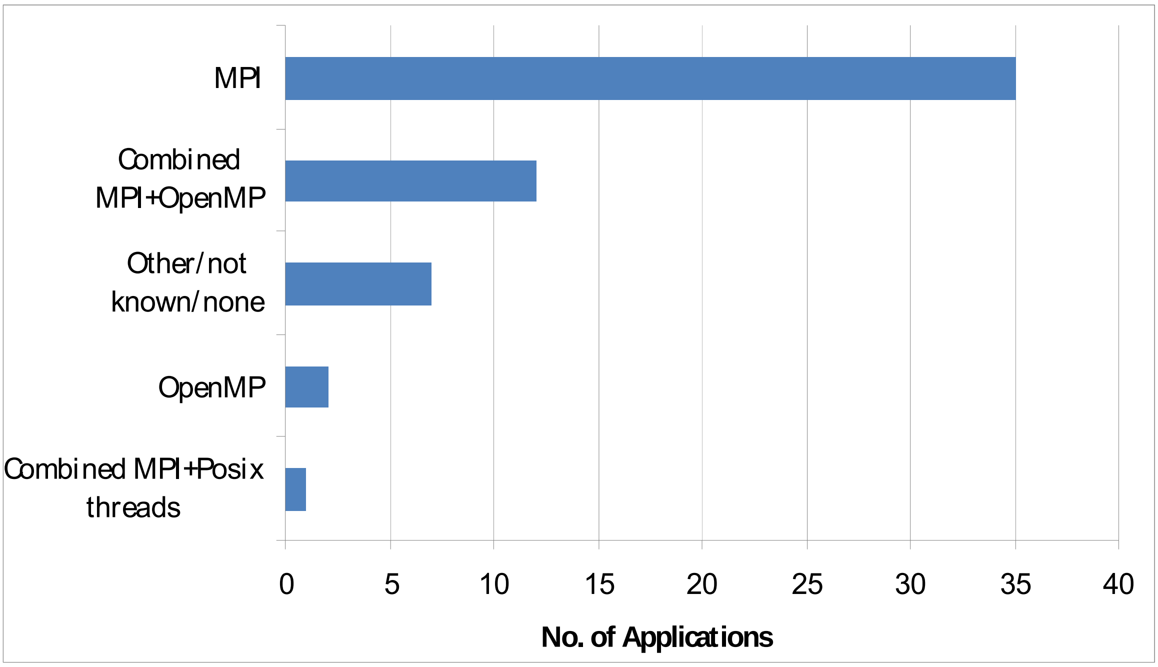

Figure 16: Parallelisation method utilisation.

Figure 16 shows the parallelisation method used for the applications. MPI is still by far the preferred method, though the proportion of applications using hybrid MPI and shared memory has increased compared to the 2008 PRACE-PP survey. No applications were reported as using HPF, nor any of the PGAS family of libraries/languages (e.g. CAF, UPC, SHMEM).

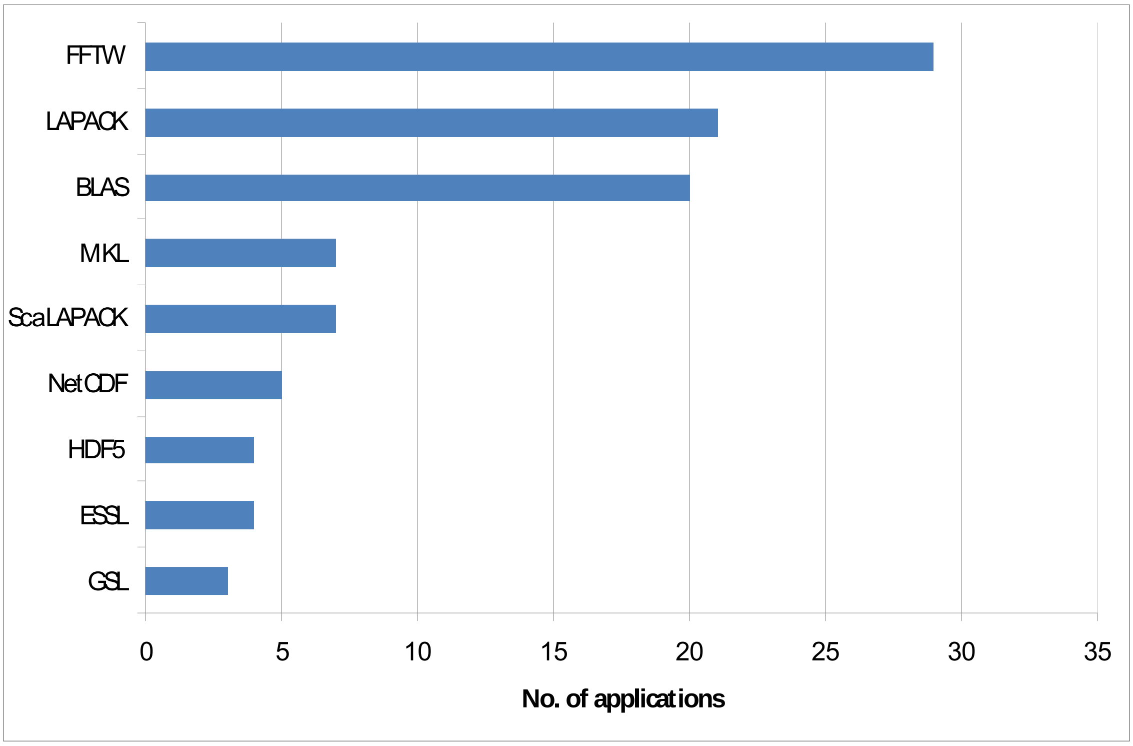

Figure 17: Scientific and I/O library utilisation.

Figure 17 shows the utilisation by the applications of scientific and I/O libraries. Not surprisingly, the most used libraries are FFTW, and the dense linear algebra family (ScaLAPACK, LAPACK and BLAS). A relatively small number of applications use I/O libraries (NetCDF and HDF5). Some applications make use of vendor specific libraries (MKL and ESSL).

3 User Survey

The user survey conducted consisted of a set of 50 questions. The respondents were largely self-selecting: any HPC users with an interest in PRACE were encouraged to complete the survey. The responses were collected between 23rd November 2010 and 17th January 2011, and a total of 411 valid responses were received. The responses to these questions are summarised below. For each question we report the response rate, i.e. the percentage of respondents who answered this particular question. We then present a summary of the responses, in graphical form wherever possible, and where appropriate comment on the results. For questions which required respondents to choose a single answer, we normally present the results as a pie-chart: where more than one answer was allowed, we use a bar chart. In some cases we have constructed additional response categories based on responses in textual format.

Question 1: Your email address &
Question 2: Your organisation / working group

Response rate: 100%

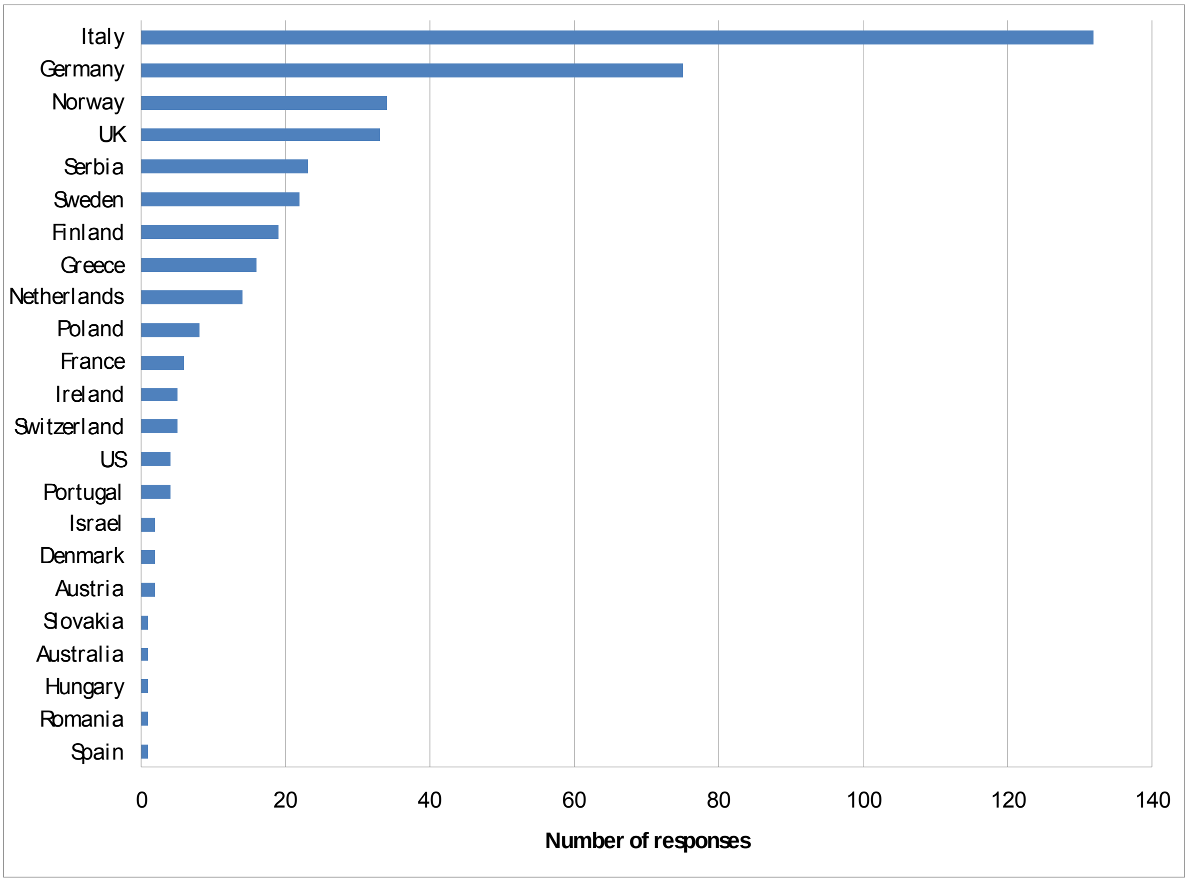

Figure 18: Summary of responses to Questions 1 and 2.

Comment: For confidentiality reasons, we have summarised these responses by nationality only.

Question 3: What scientific field are you working in? Please try to select from the following scientific areas, only use *other* if absolutely necessary.

Response rate: 100%

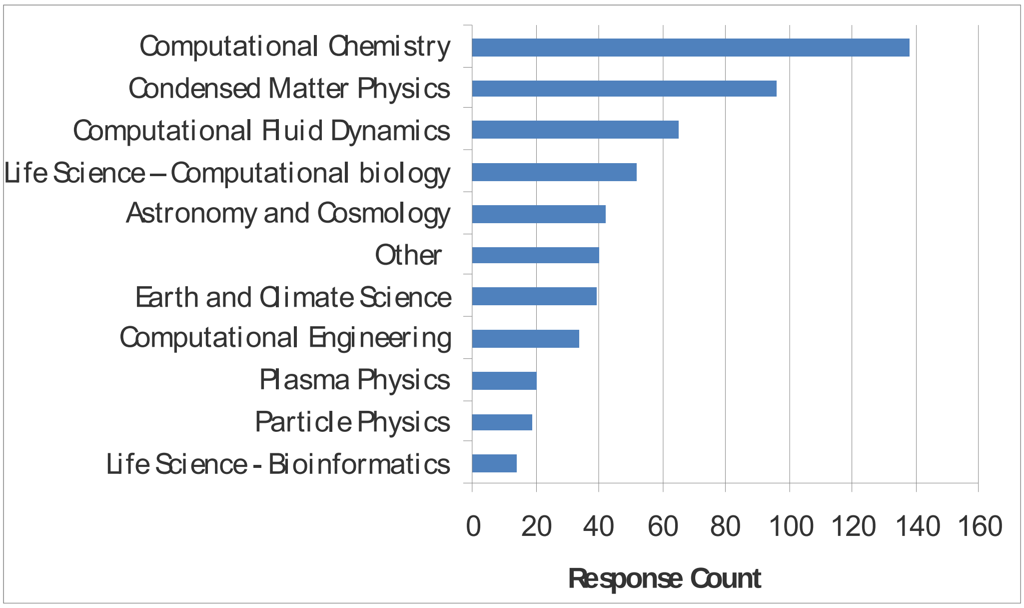

Figure 19: Summary of responses to Question 3.

Comment: The responses in the "Other" category include various fields in Informatics and Numerical Mathematics.

Question 4: Which machine are you mainly working on? If you are working on more than one system, please select the major system that you are working on.
&
Question 5: Which organisation hosts this machine?

Response rate: 100%

Comment: Because the answers to this question were in free text format, and many respondents did not name a specific system, it has not been possible to provide a reliable summary of responses.

Question 6: What is the main application you are using on this system?

Response rate: 100%

Application	*Response Count*
NAMD	18
Gaussian	17
VASP	16
Gromacs	14
CRYSTAL	9
Own Code	9
Quantum Espresso	9
Molecular dynamics	8
CP2K	7
CPMD	7
GAMESS	7
Lattice QCD	6
Dalton	5
GADGET	4
Monte Carlo simulation	4
OpenFOAM	4
CFD	3
Continuous Time Quantum Monte Carlo	3
NWChem	3
PLUTO	3
SIESTA	3
Yambo	3
ABINIT	2
BQCD	2
CARP	2
Chaste	2
Chroma	2
Desmond	2
NEURON	2
Pencil Code	2
SPEEDUP	2
TAU	2
TurboRVB	2
Wien2k	2
Total	186

Table 5: Summary of responses to Question 6.

Comment: Table 5 lists applications which were recorded by more than one user. The remaining responses either listed a code which was unique to that user, or else mentioned only a general scientific area.

Question 7: What is your relationship with the code?

Response rate: 100%

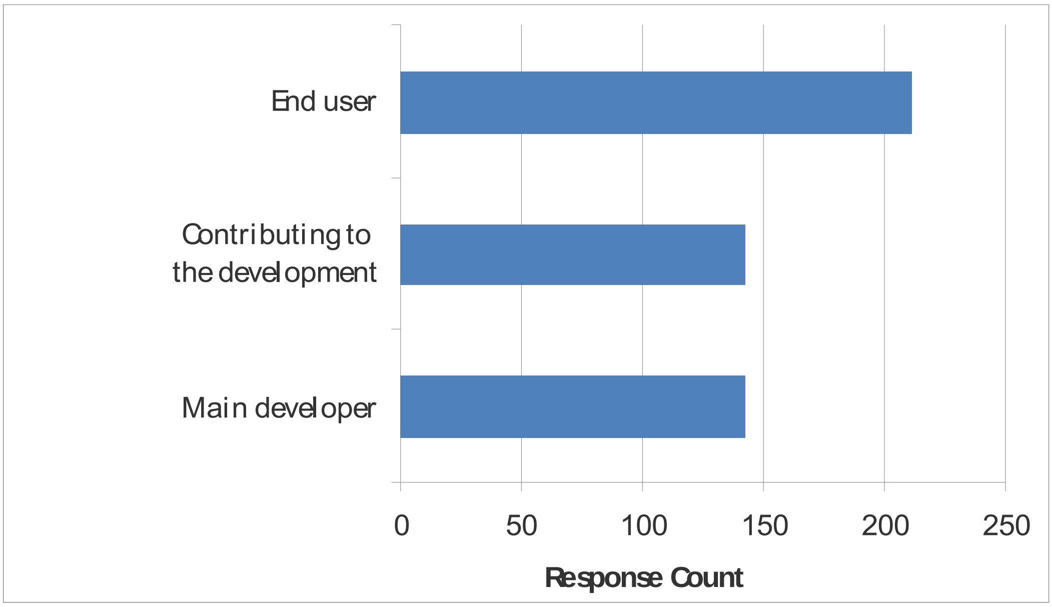

Figure 20: Summary of responses to Question 7.

Comment: It is interesting to note that the responses are fairly evenly balanced between the three categories.

Question 8: What is the licensing arrangement for the application?

Response rate: 96%

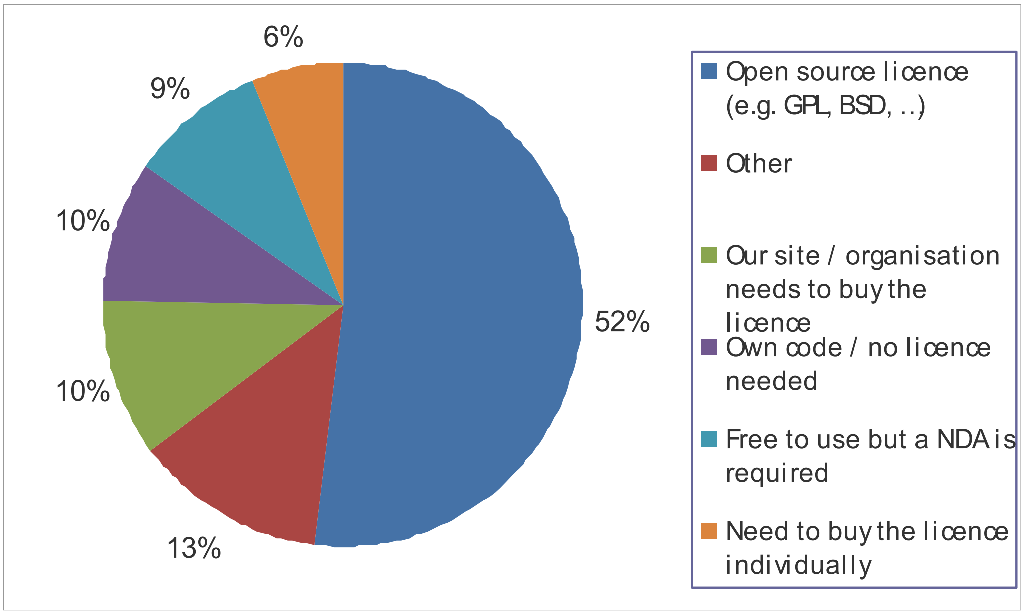

Figure 21: Summary of responses to Question 8.

Question 9: What is the range of applicability of the code?

Response rate: 99%

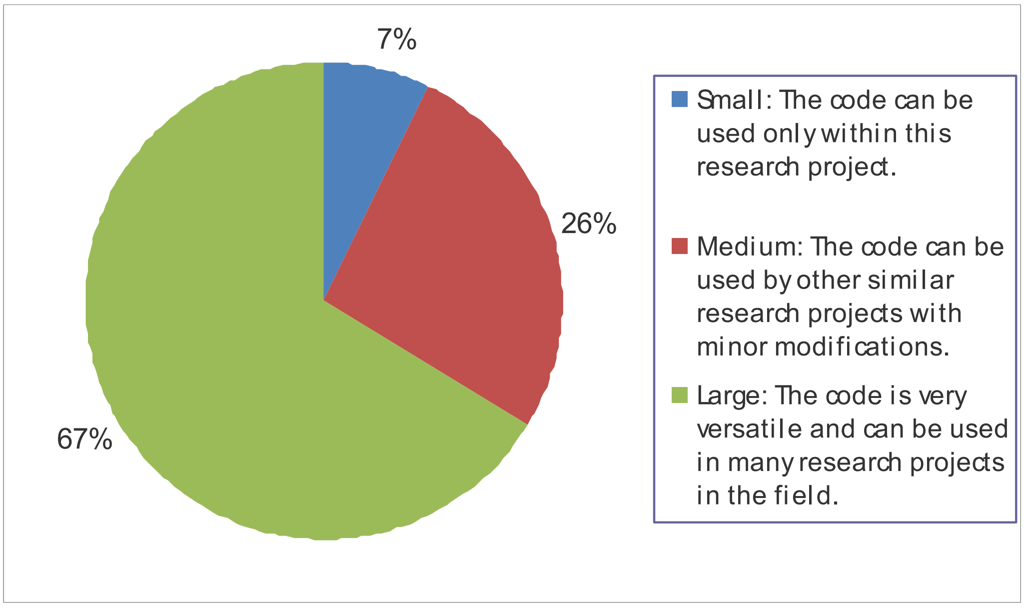

Figure 22: Summary of responses to Question 9.

Question 10: Who ported the code to the system for your research?

Response rate: 98%

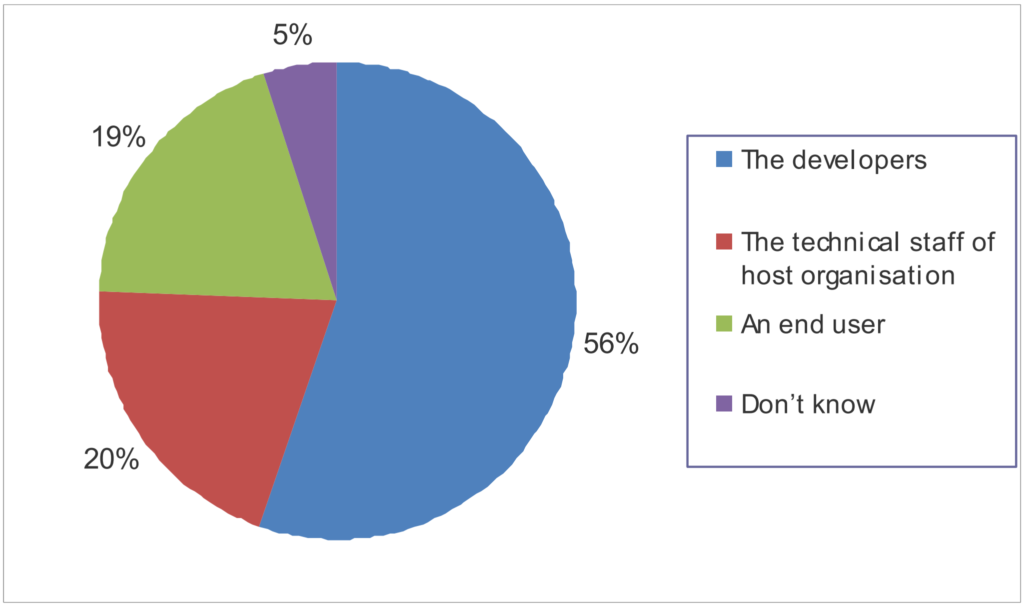

Figure 23: Summary of responses to Question 10.

Question 11: Please select from the following list: what is the current scalability of the production job for you application on this machine and what is the desired / required scalability after the application enabling work?

Response rate: 97%

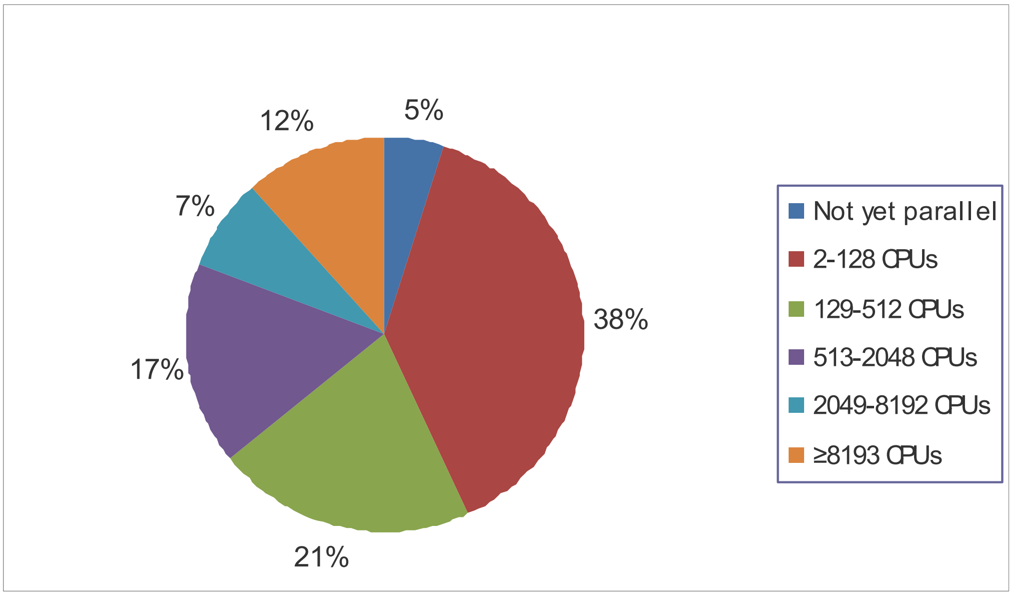

Figure 24: Summary of responses to Question 11: Current scalability.

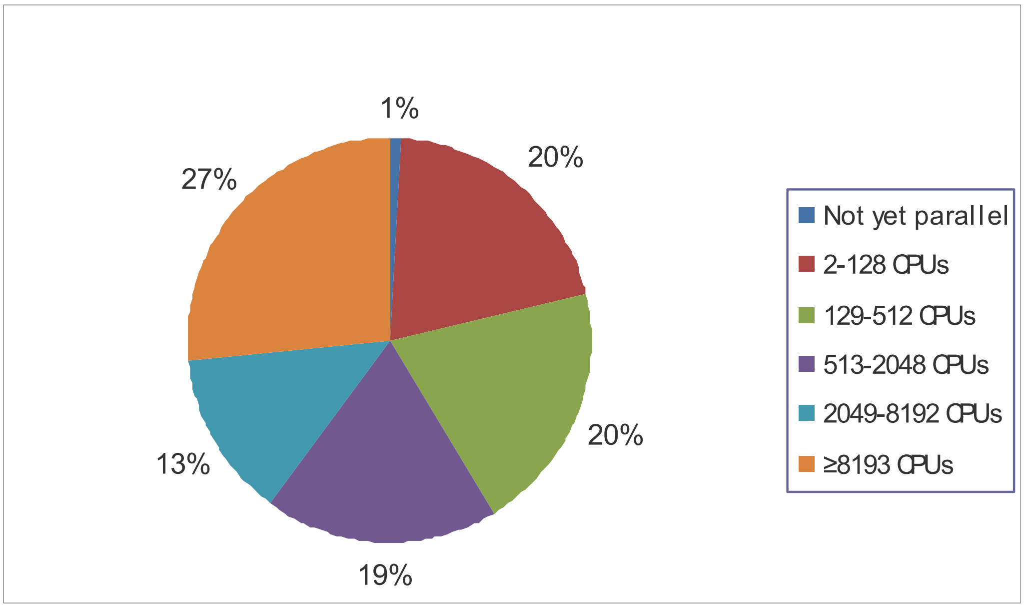

Figure 25: Summary of responses to Question 11: Desired/required scalability

Comment: Figure 24 shows the responses for current scalability, while Figure 25 shows those for desired/required scalability**.** These results show a clear desire from most users to increase the scalability of the applications they are using. On average, however, the desired increase in scalability is relatively modest: less than an order of magnitude in the number of CPUs.

Question 12: On which type of scalability do your computations rely?

Response rate: 92%

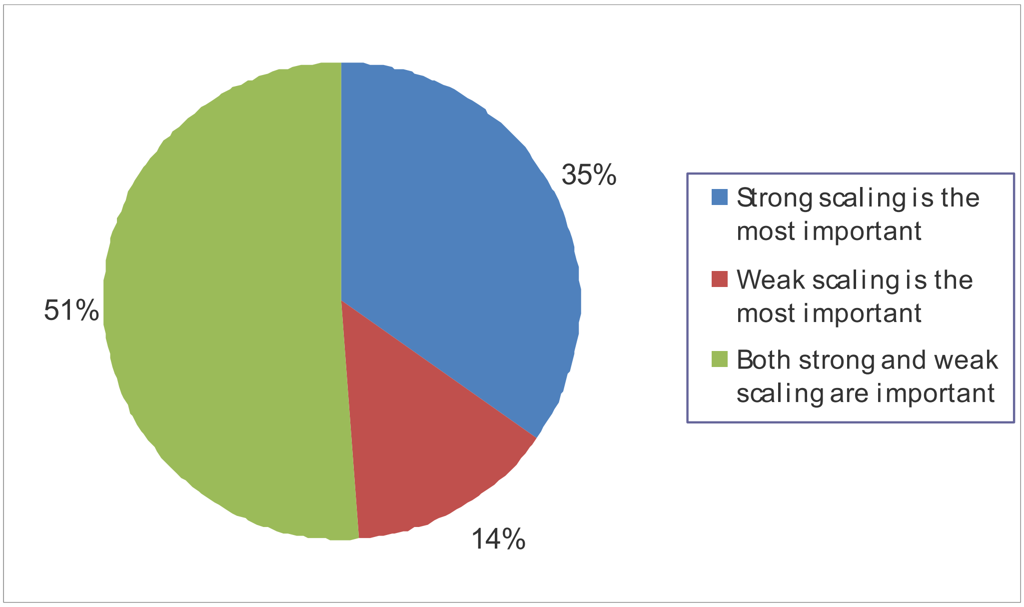

Figure 26: Summary of responses to Question 12.

Comment: These results show that both strong scaling (where the overall problem size is kept constant while increasing the number of cores used), and weak scaling (where the overall problem size increases in proportion to the number of cores used) are important to a significant number of users.

Question 13: What memory size per core is required for your production job?

Response rate: 99%

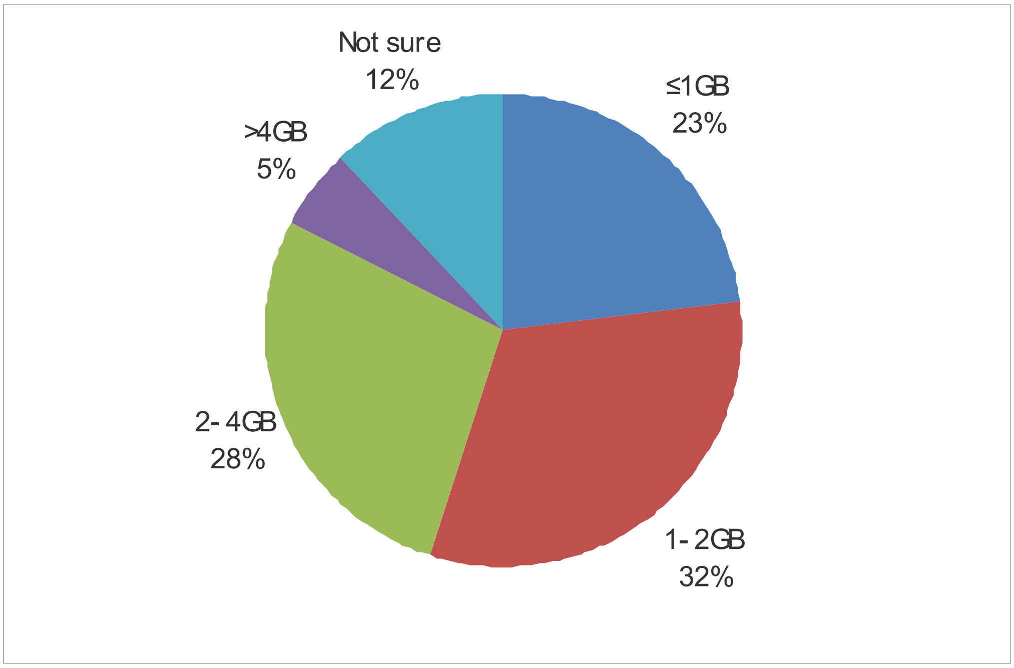

Figure 27: Summary of responses to Question 13.

Comment: If architectures continue to evolve to have more cores per node, and less memory per core, memory requirements may become a significant obstacle for a large number of users.

Question 14: What are your expectations with regards to the performance of your application on this machine?

Response rate: 96%

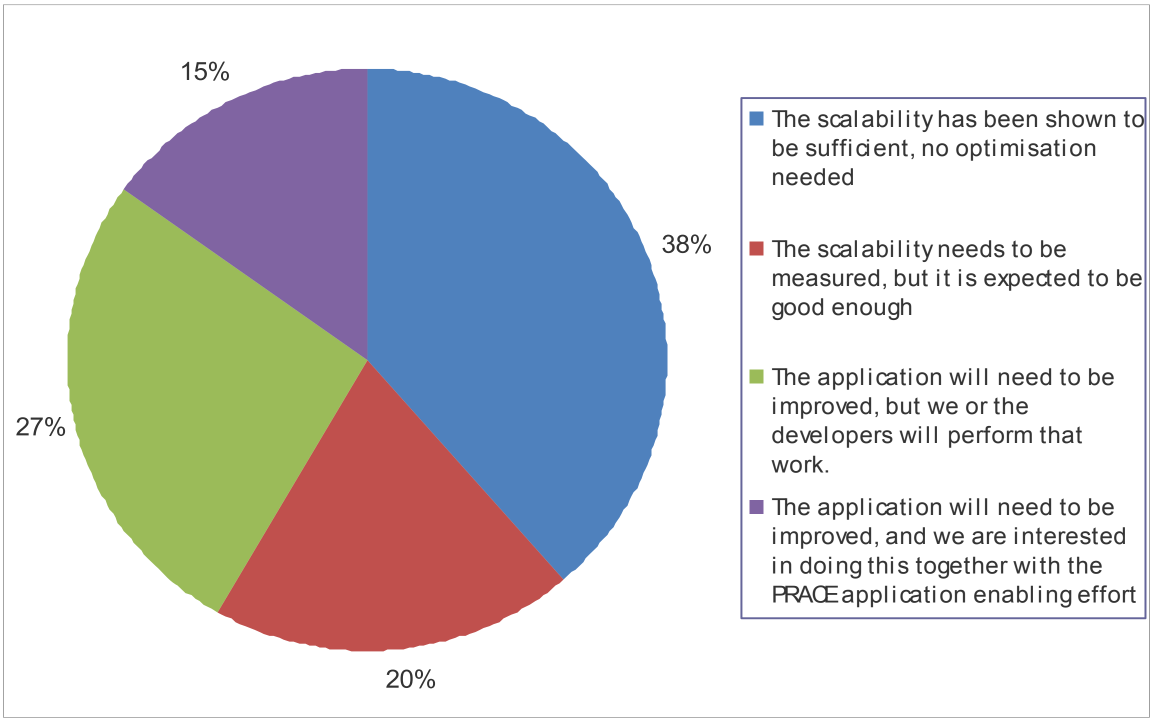

Figure 28: Summary of responses to Question 14.

Question 15: Do you have a clear grasp of what are the performance limiting factors in your code?

Response rate: 96%

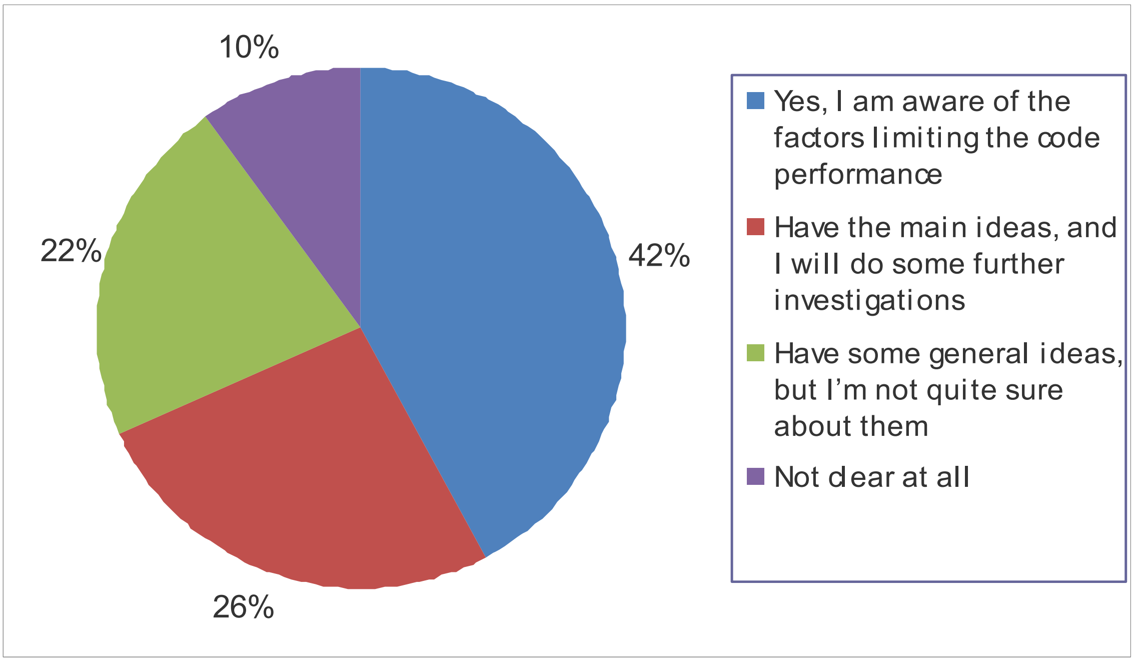

Figure 29: Summary of responses to Question 15.

Question 16: How much effort is needed to enable the code to scale to larger systems?

Response rate: 95%

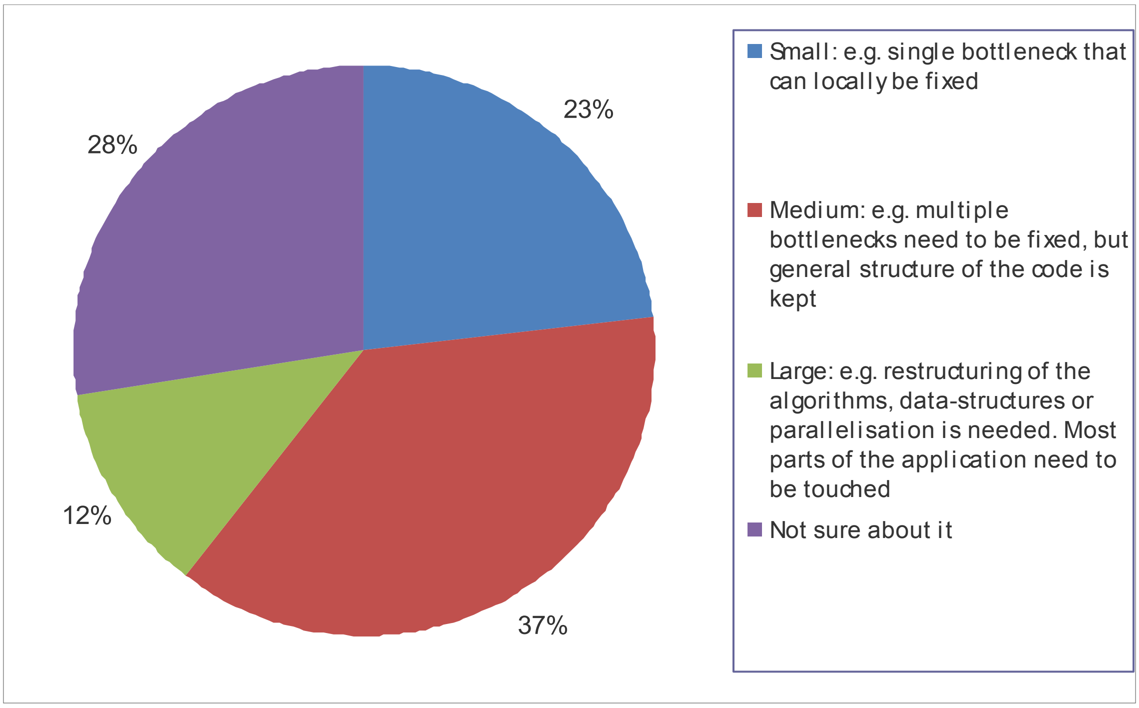

Figure 30: Summary of responses to Question 16.

Question 17: Please describe in detail what the current challenges / issues are met for your application scaling / performance improvement.

Response rate: 41%

Comment: This question allowed a free format of responses. There was a wide variety of responses, but recurring themes included the need for new algorithms; communication, memory bandwidth and I/O bottlenecks; and the need to exploit hybrid (MPI plus shared memory) programming.

Question 18: Would you consider migrating to an alternative, more scalable or more efficient, production code if necessary in order to access the PRACE Tier-0 systems?

Response rate: 94%

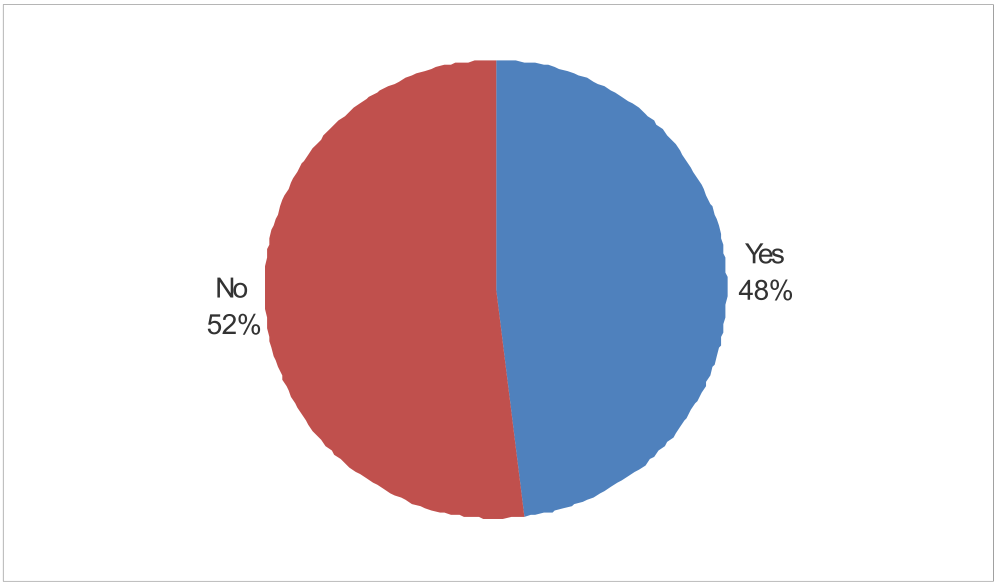

Figure 31: Summary of responses to Question 18.

Comment: The response to this question shows a perhaps surprisingly high degree of willingness for scientists to consider alternative, more scalable, codes.

Question 19: To try and assess the requirements for the PRACE systems quantitatively, we would like you score the following architecture features of hardware in terms of importance of your code. A total of 20 points should be distributed amongst the following requirements, with higher priority requirements receiving a higher number of points. If features are not important at all a score of zero can be used.

Response rate: 62%

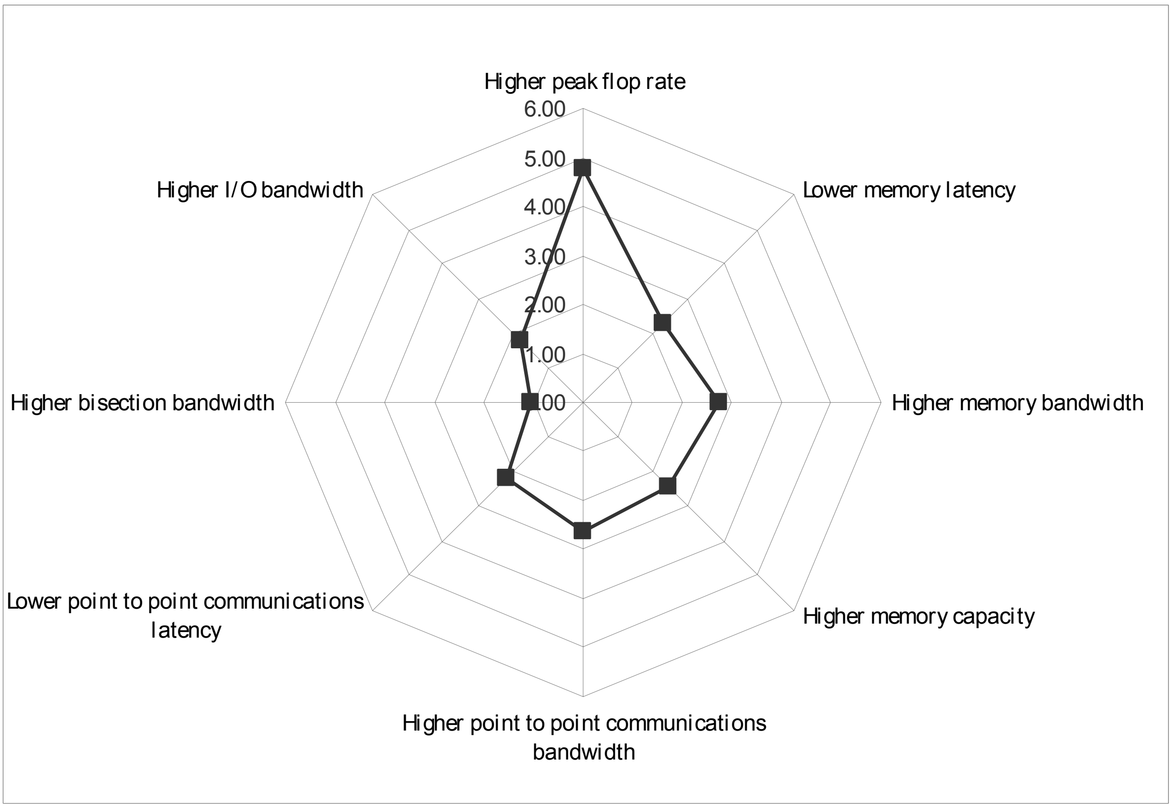

Figure 32: Summary of responses to Question 19.

Comment: Of the listed features, all were considered important. A higher peak flop rate scored higher than the other features, while higher bisection bandwidth (the minimum bandwidth between two halves of the system, considering all possible ways to split the nodes into two equal partitions) scored lower. The other features all received similar scores on average.

Question 20: Are there any other architectural features that might affect the performance of your application?

Response rate: 19%

Comment: The number of relevant responses to this question was low. The following features were mentioned by a small number of users:

- larger cache sizes
- vector/SIMD capability
- GPGPUs.

Question 21: Please select one or more from the following sentences which describe the mechanism your application adopts to access data

Response rate: 73%

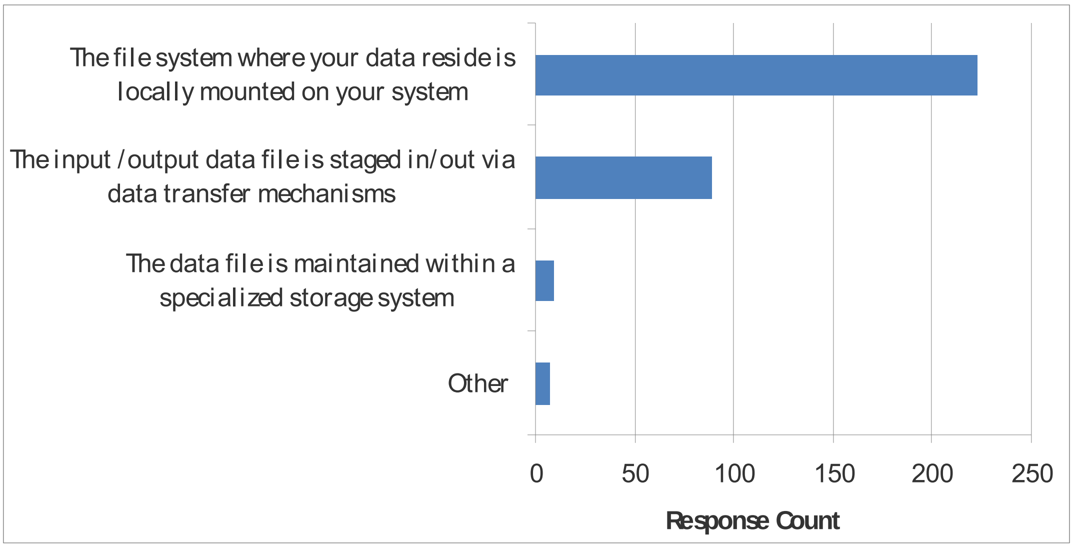

Figure 33: Summary of responses to Question 21.

Comment: For the majority of users, data resides on the local file system, while a substantial minority (almost 30%) stages their data in/out. Only a few use specialized storage systems.

Question 22: Concerning the data transfer between other machines and this system (data staging-in and out), which of the following is more important for your application?

Response rate: 73%

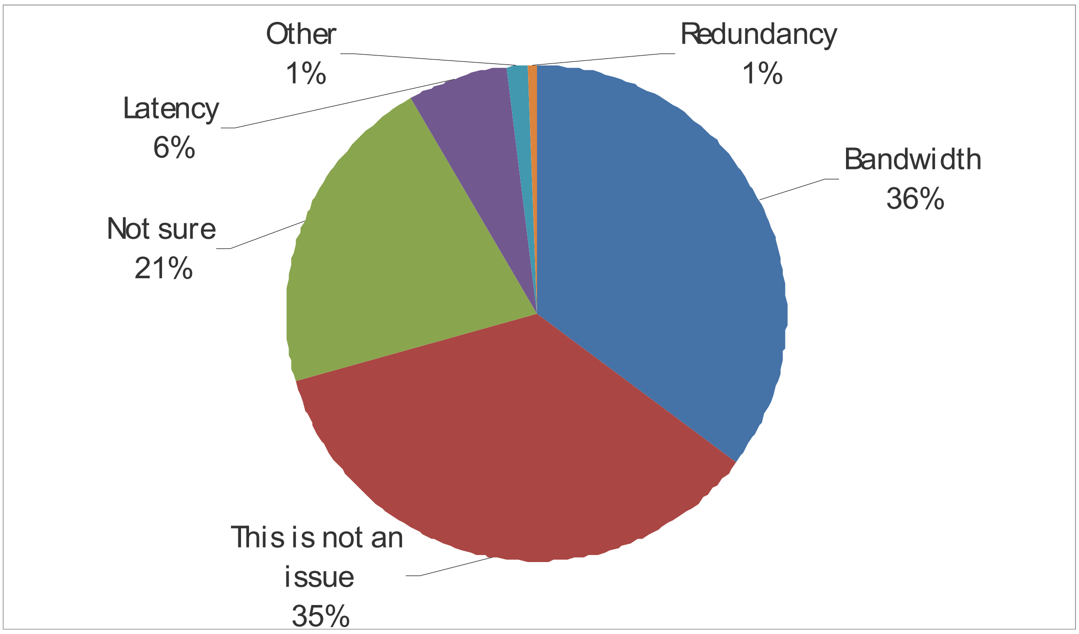

Figure 34: Summary of responses to Question 22.

Question 23: What is the minimum amount of disk space required per production job (in GB)?

Response rate: 72%

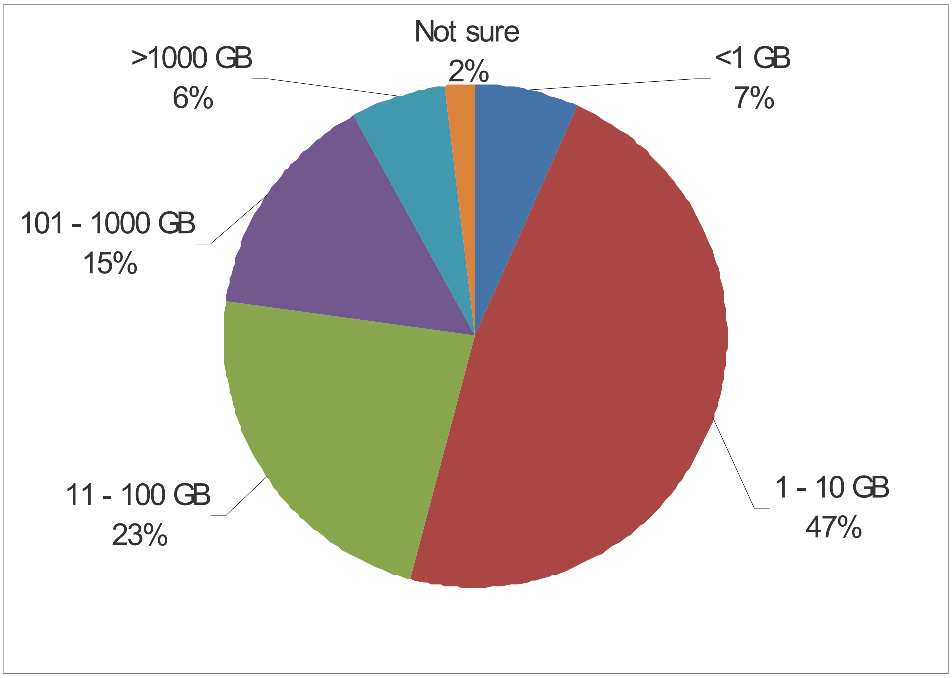

Figure 35: Summary of responses to Question 23.

Comment: There is very wide range of requirements reported here, spanning more than three orders of magnitude. It should be noted, however, that users may have different notions of what constitutes a production job.

Question 24: On a monthly basis, by how much would you expect your overall data storage usage to increase (in GB)?

Response rate: 67%

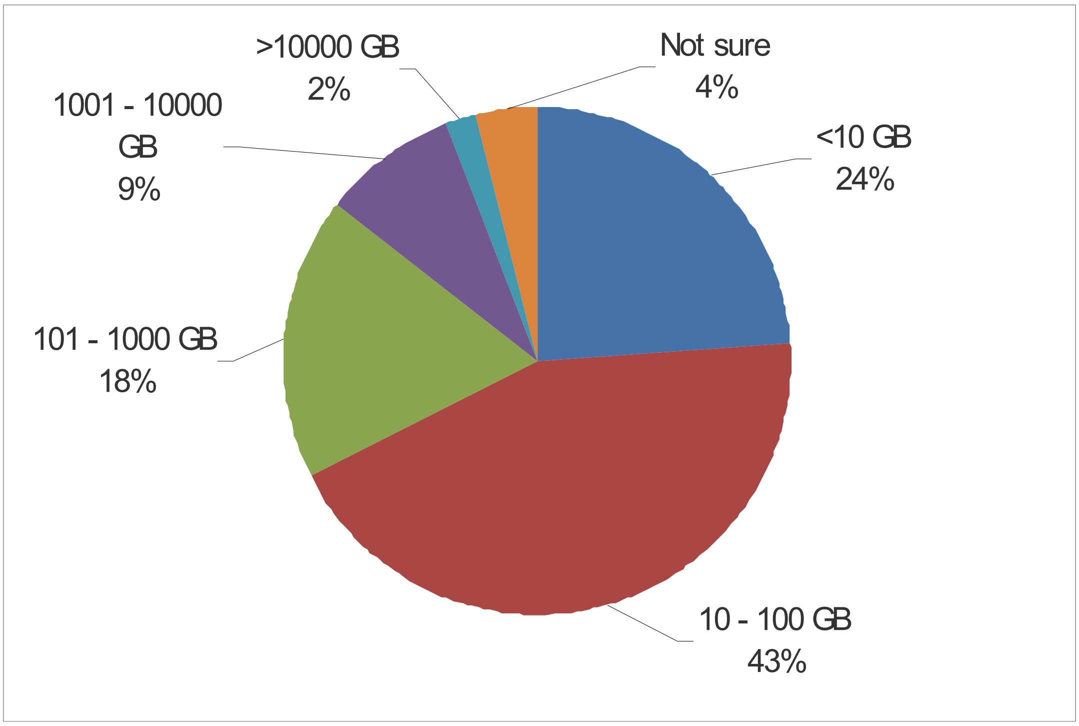

Figure 36: Summary of responses to Question 24.

Comment: As for Question 23, there is very wide range of requirements reported here.

Question 25: On a monthly basis, how much data would you expect to transfer on/off the system (in GB)?

Response rate: 67%

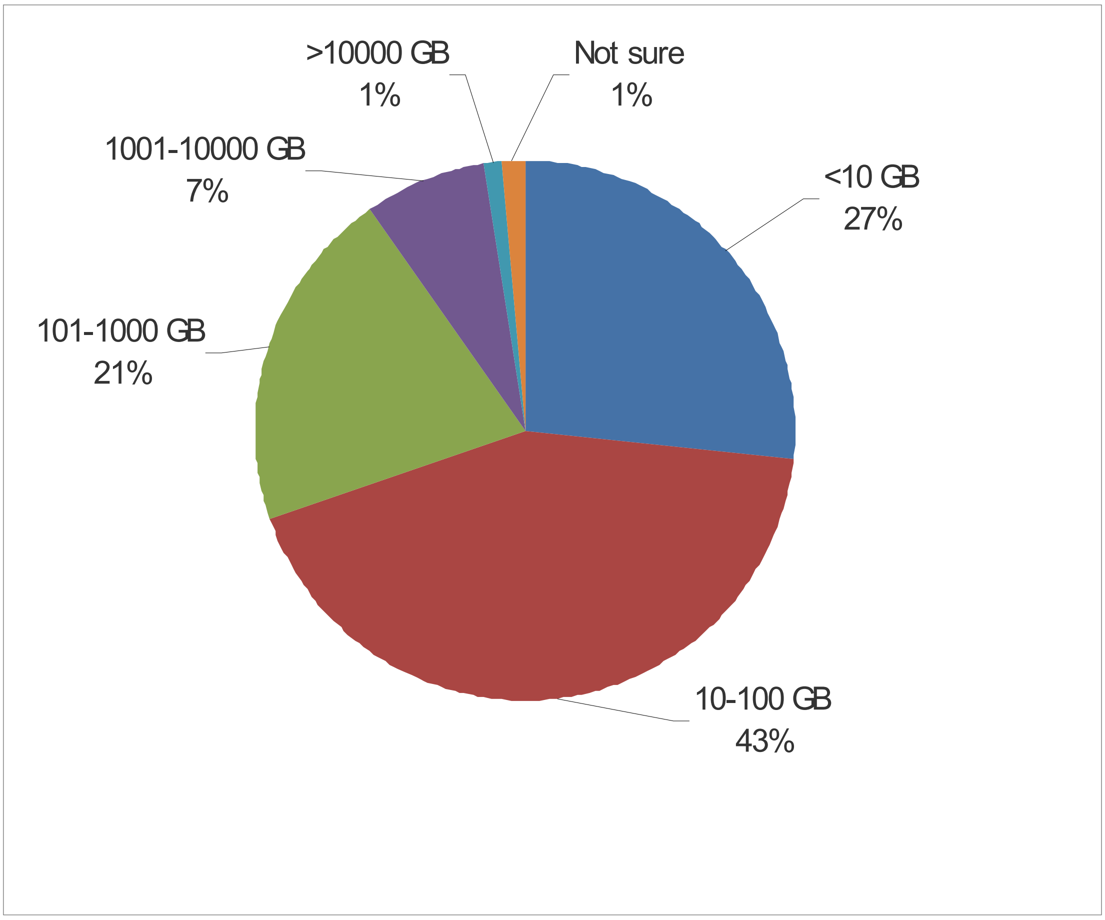

Figure 37: Summary of responses to Question 25.

Comment: The responses to this question are very similar to those for Question 24.

Question 26: What is the typical size of your individual data files (in GB)?

Response rate: 70%

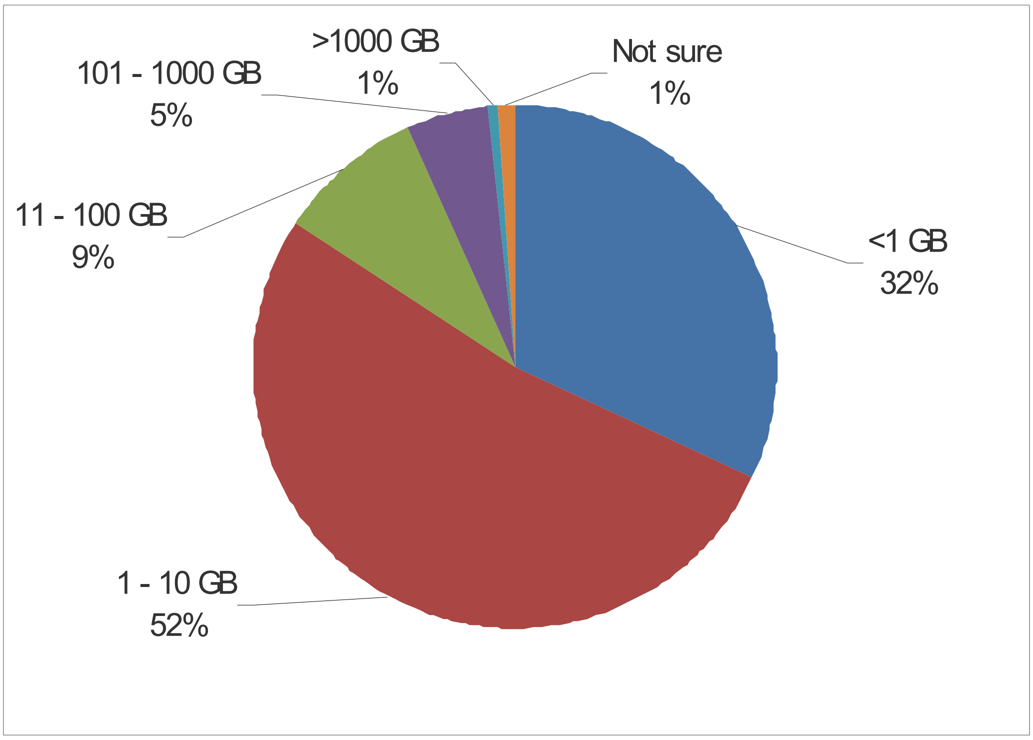

Figure 38: Summary of responses to Question 26.

Comment: Again, there is a wide range of files sizes used, though for 85% of users the typical file size is less than 10GB.

Question 27: Do you have any requirements to share data with other users (such as one-off transfers, or commonly held files)? Please give details.

Response rate: 28%

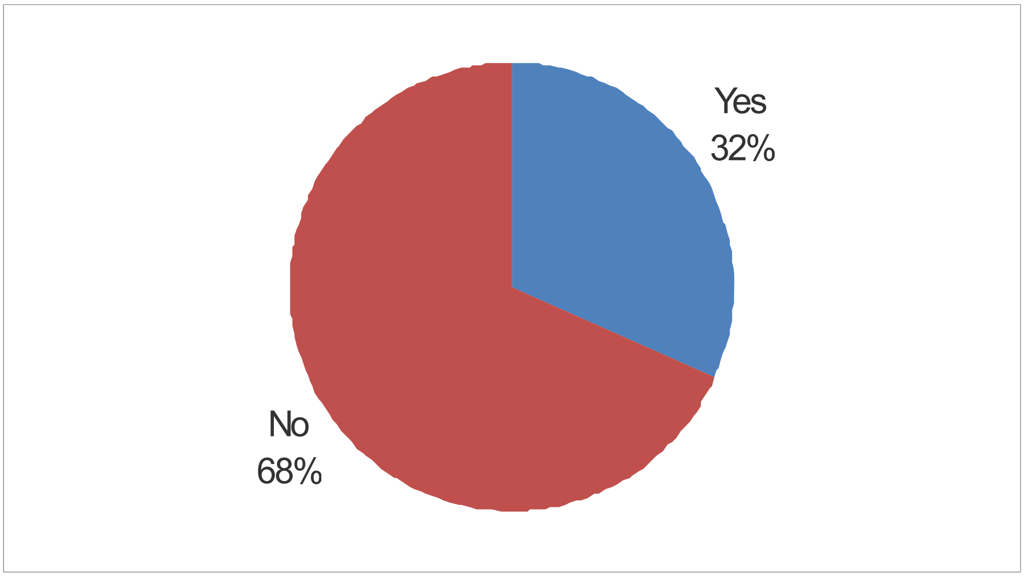

Figure 39: Summary of responses to Question 27.

Comment: Many of the "yes" responses referred to the need for shared file space on the production system. Note the low response rate: it may be reasonable to assume that the 72% of users who skipped this question also do not have any such requirements.

Question 28: Which of the following data transfer mechanisms are you familiar with?

Response rate: 73%

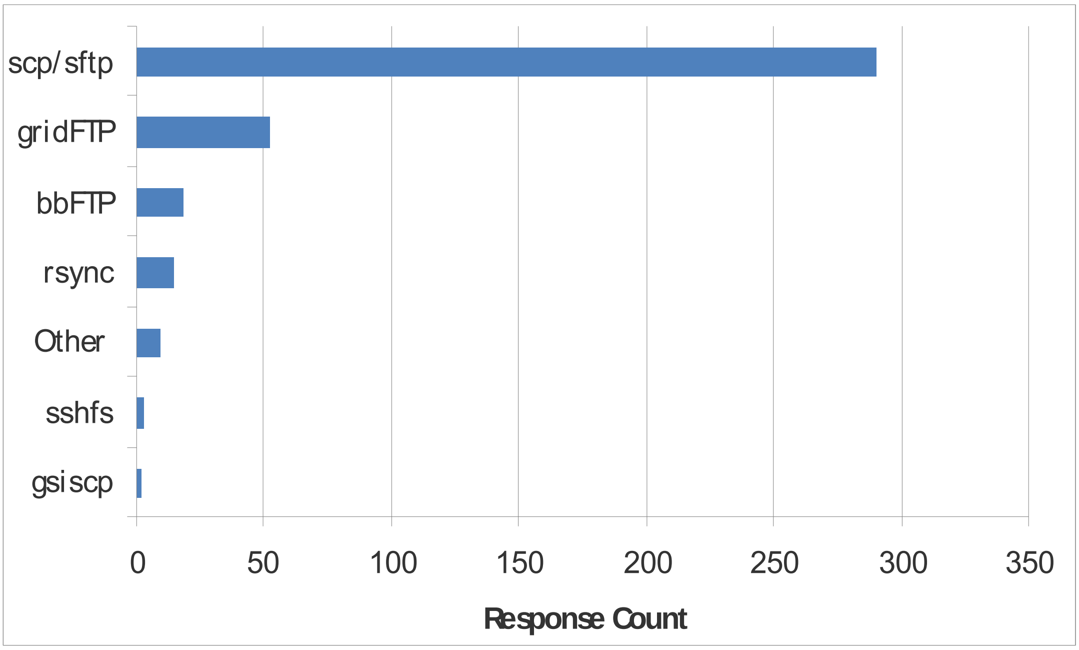

Figure 40: Summary of responses to Question 28.

Question 29: Do you have any requirements for network connectivity for running applications (e.g. Meta-Computing, Computational Steering, Real-time Visualisation)?

Response rate: 70%

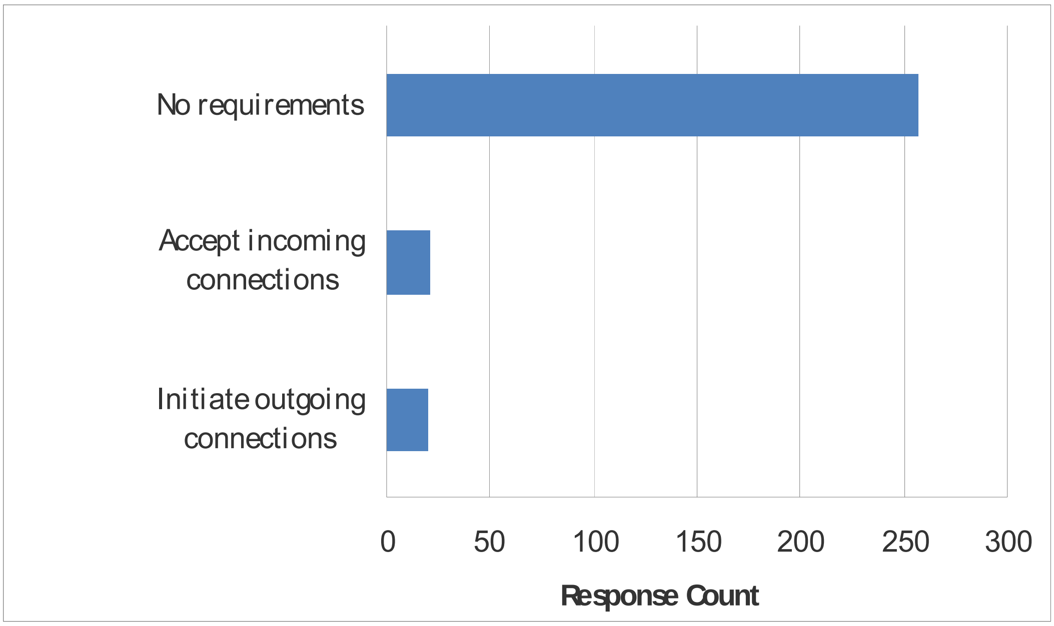

Figure 41: Summary of responses to Question 29.

Question 30: Relating to Question 29 above, please list the required software if possible.

Response rate: 3%

Comment: A very small number of users mentioned software by name. These included VisIt, ParaView, MPWide, VMD and CHIMERA.

Question 31: Which programming models and languages do you use for code development? Please select one or more from the following list.

Response rate: 78%

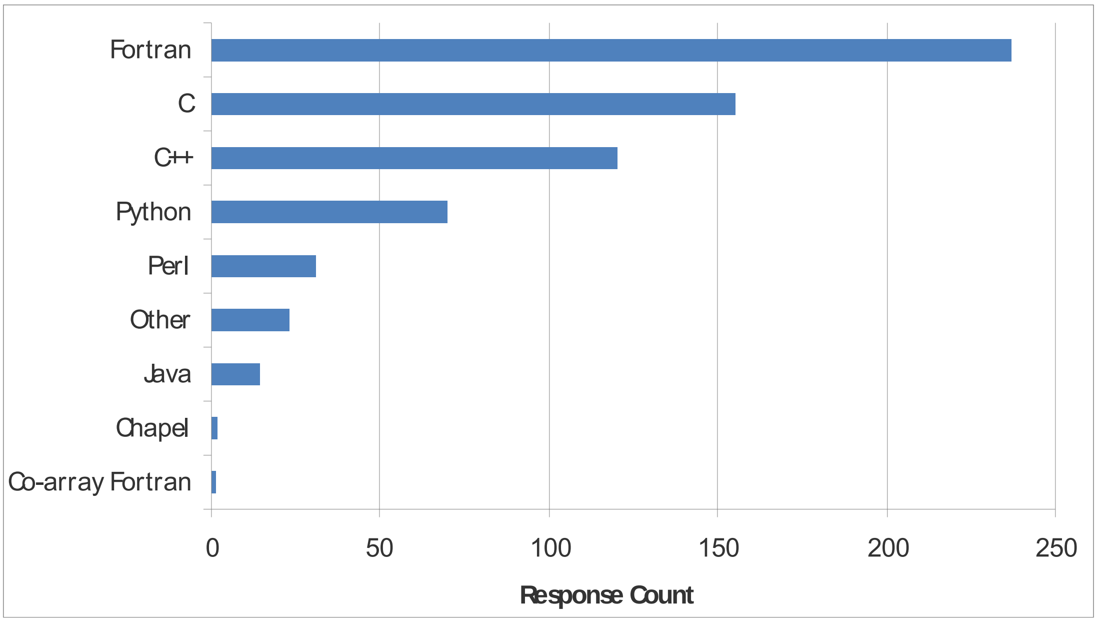

Figure 42: Summary of responses to Question 31.

Comment: Languages mentioned by more than one user in the “other” category were Matlab, R, CUDA, Unix shell script and awk. There was a zero response for UPC and Fortress. It is interesting to note that the sum of responses for C and C++ outnumber those for Fortran.

Question 32: Which parallelisation implementations do you use for code development? Please select one or more from the following list.

Response rate: 72%

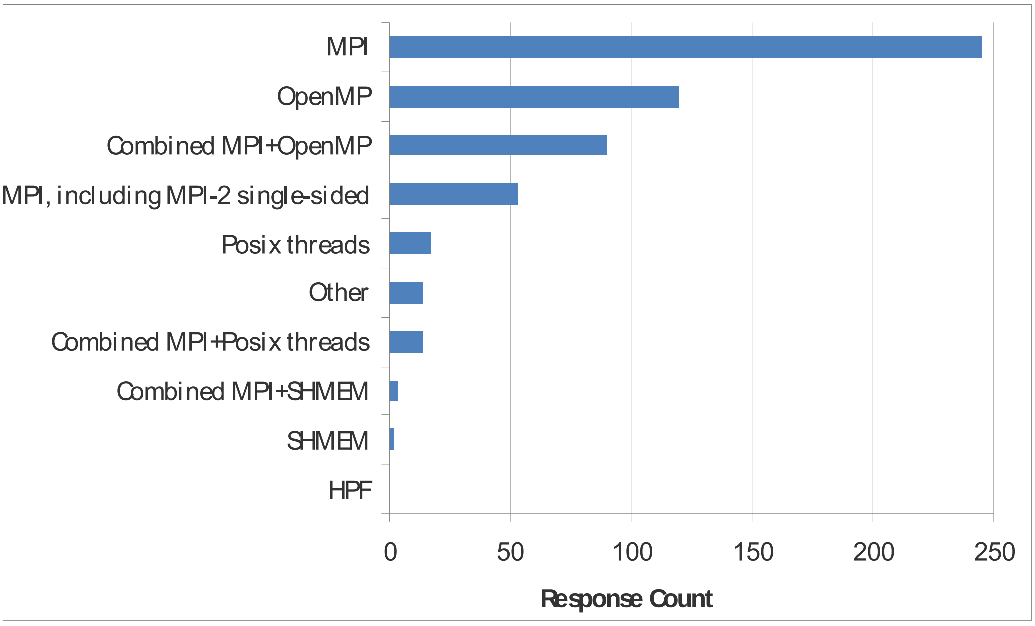

Figure 43: Summary of responses to Question 32.

Comment: The response count for MPI-2 single sided was surprisingly high: it is possible that some users chose this category, even if they only used MPI-1 features.

Question 33: Which of the following batch-schedulers are you familiar with?

Response rate: 70%

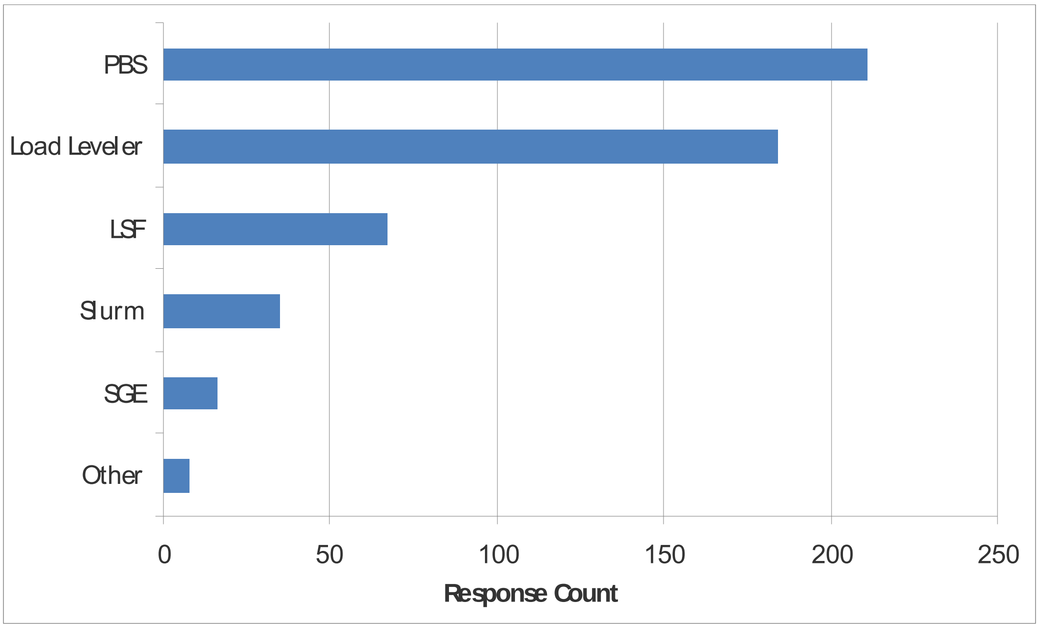

Figure 44: Summary of responses to Question 33.

Question 34: Is HPC time that is allocated to your research group also used by collaborators that are based in another country?

Response rate: 76%

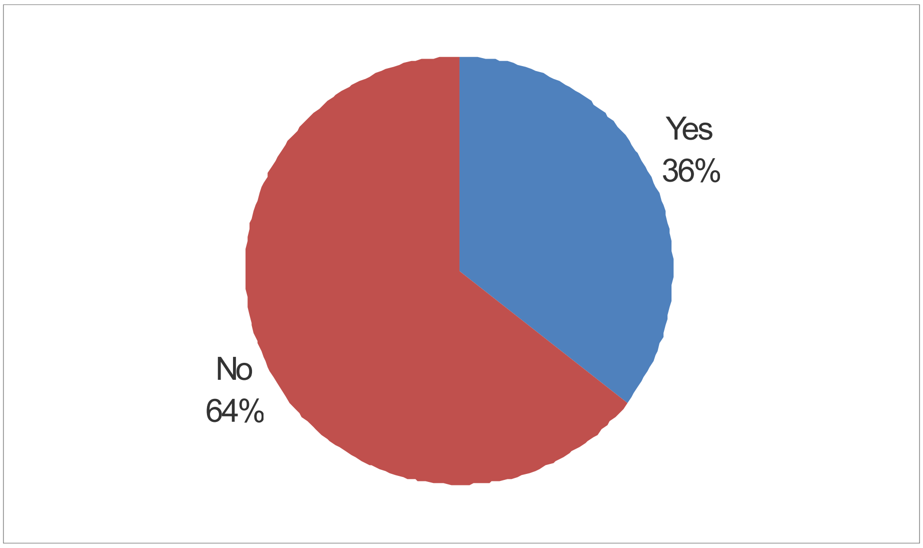

Figure 45: Summary of responses to Question 34.

Question 35: Are you familiar with X.509 certificate-based authentication? Do you hold a X.509 certificate issued by a European trusted CA?

Response rate: 76%

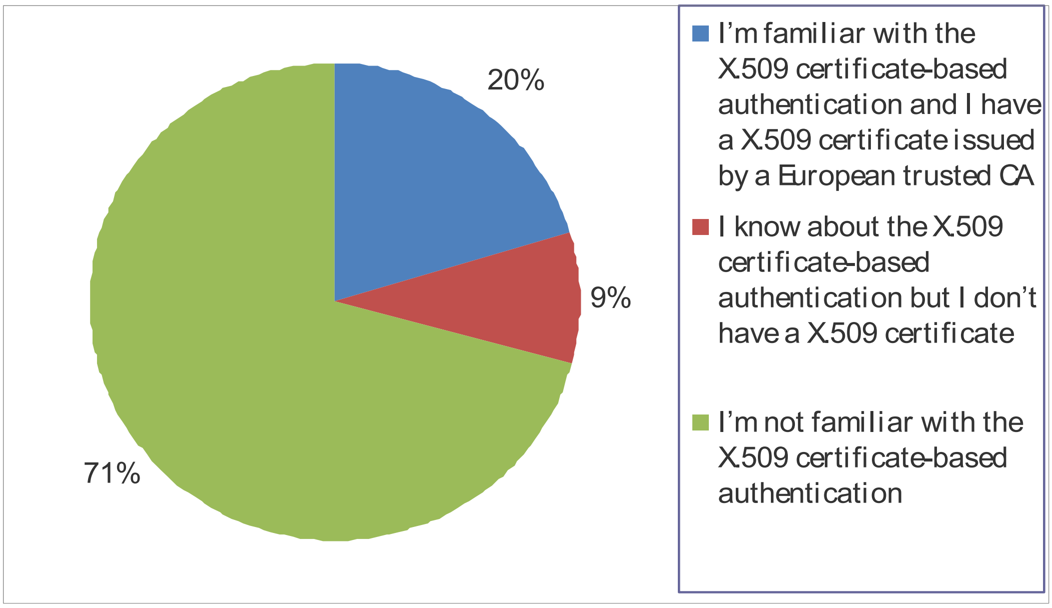

Figure 46: Summary of responses to Question 35.

Question 36: Which of the following grid middleware are you familiar with?

Response rate: 74%

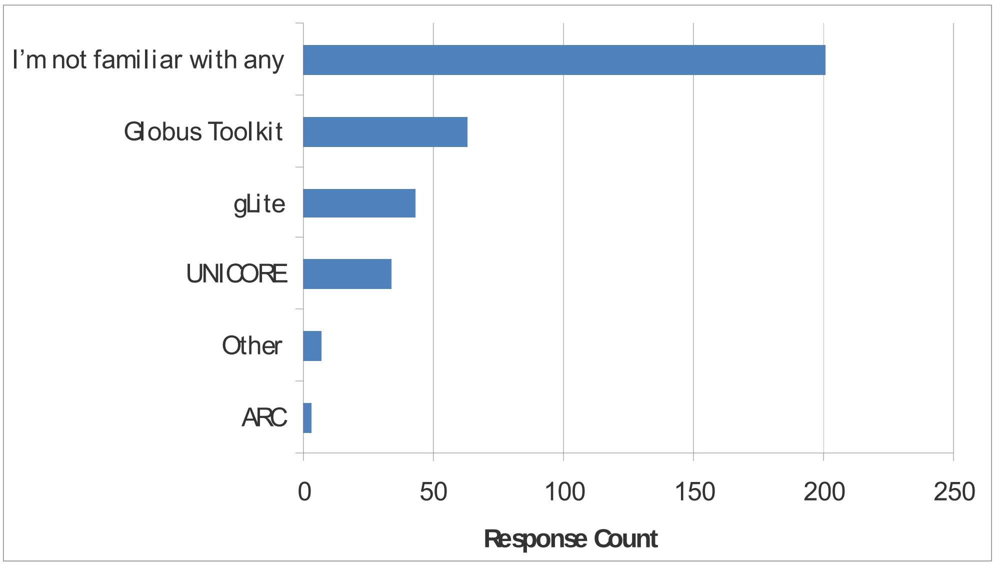

Figure 47: Summary of responses to Question 36.

Question 37: Does your application require a workflow system to be executed? Which of the following systems are you familiar with?

Response rate: 65%

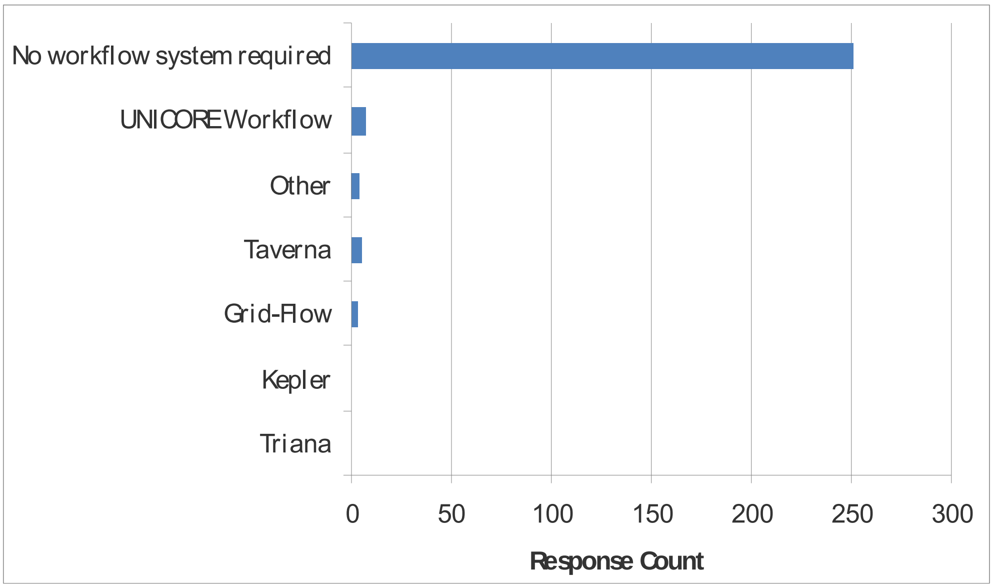

Figure 48: Summary of responses to Question 37.

Question 38: Do you need any specific tool for the visualization of your application results?

Response rate: 44%

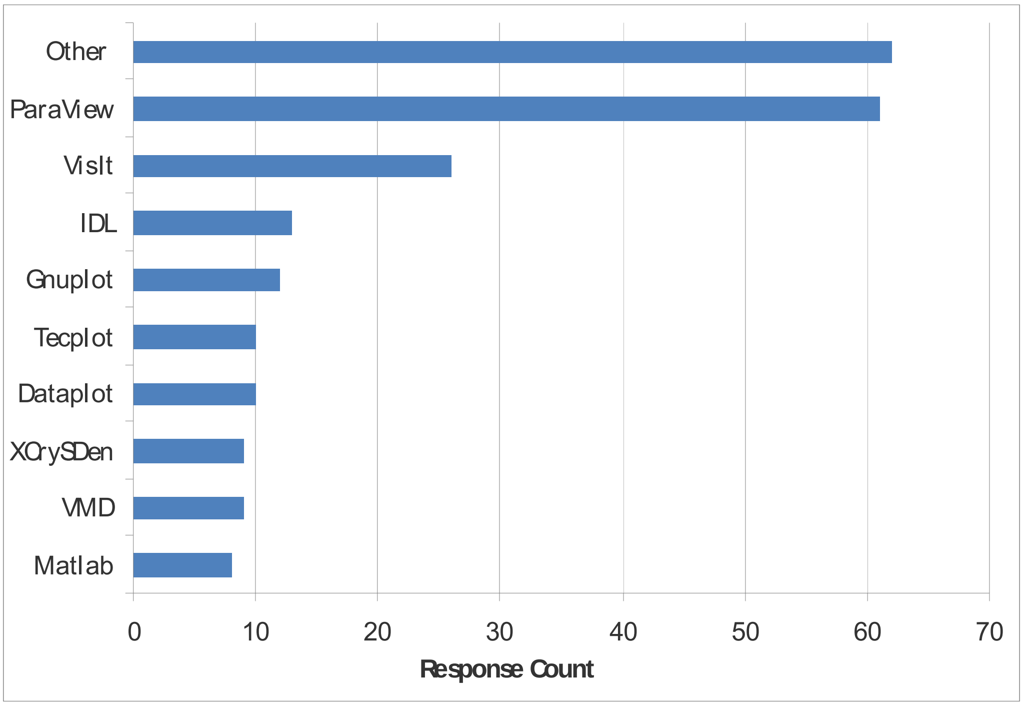

Figure 49: Summary of responses to Question 38.

Comment: A number of responses in the "Other" category mentioned specific tools developed for individual applications.

Question 39: Have you ever used any of the following ways to contact the site operator / host organisation in case of the problems with unexpected application behavior?

Response rate: 44%

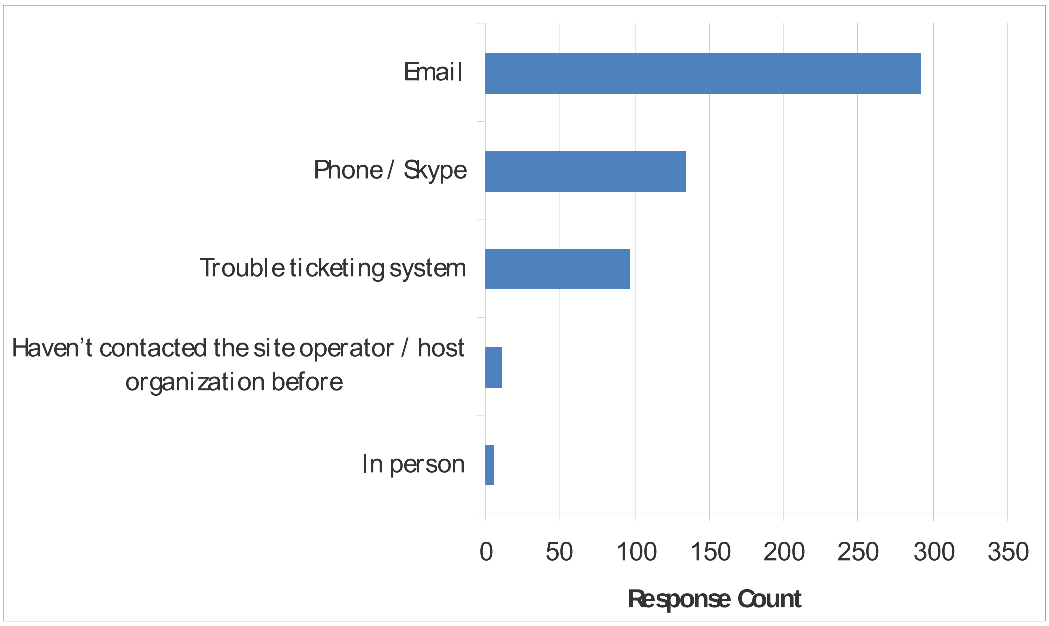

Figure 50: Summary of responses to Question 39.

Question 40: What status information of your account (e.g. resource utilisation, disk quota, etc.) are you interested in?

Response rate: 59%

Comment: Almost all respondents simply mentioned resource/CPU utilization and/or disk quota.

Question 41: Relating to question 40 above, if not available in real time, how frequently would you require such information to be updated?

Response rate: 68%

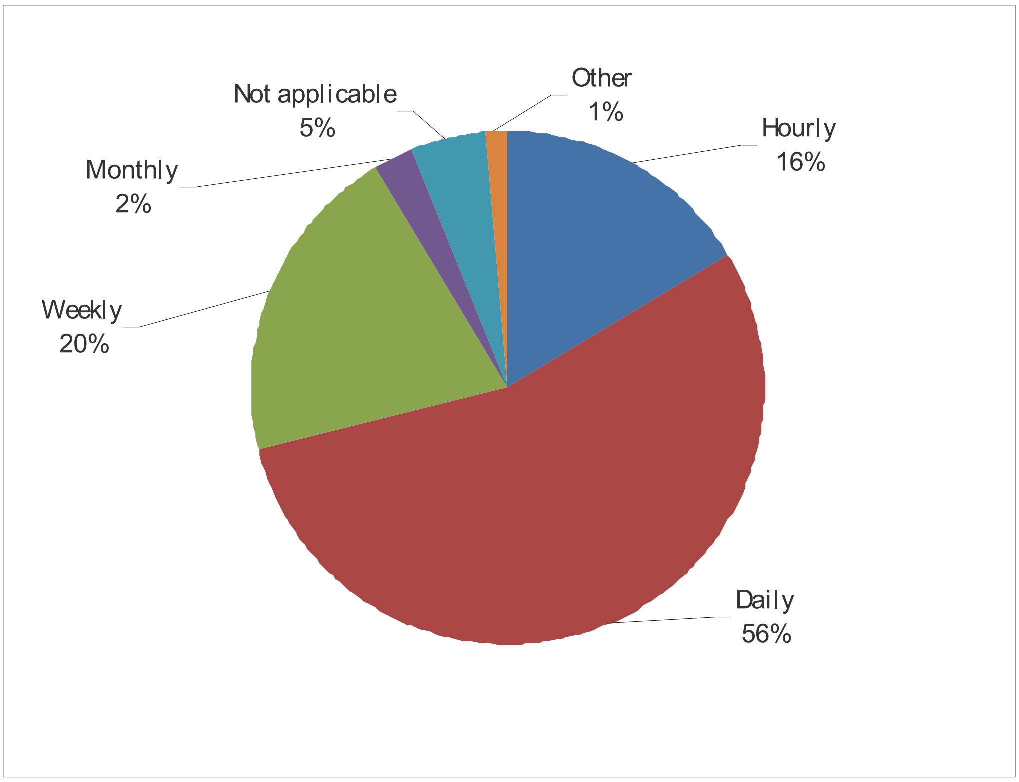

Figure 51: Summary of responses to Question 41.

Question 42: What collaborative tools (code versioning system, document repository, wiki, forum, etc.) will help you working with PRACE?

Response rate: 41%

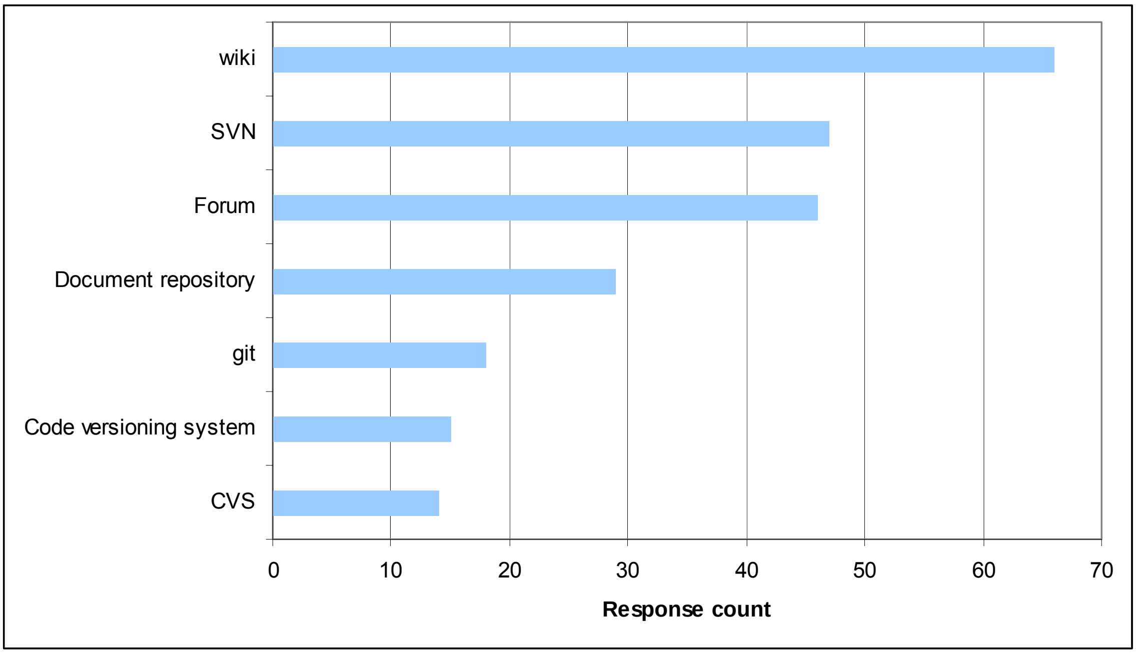

Figure 52: Summary of responses to Question 42.

Question 43: If you encountered any security issues on your site/machine account, please describe briefly here.

Comment: Responses to this question are omitted for confidentiality reasons, and are for internal project use only.

Question 44: Have you ever attended any PRACE/DEISA training events?

Response rate: 24%

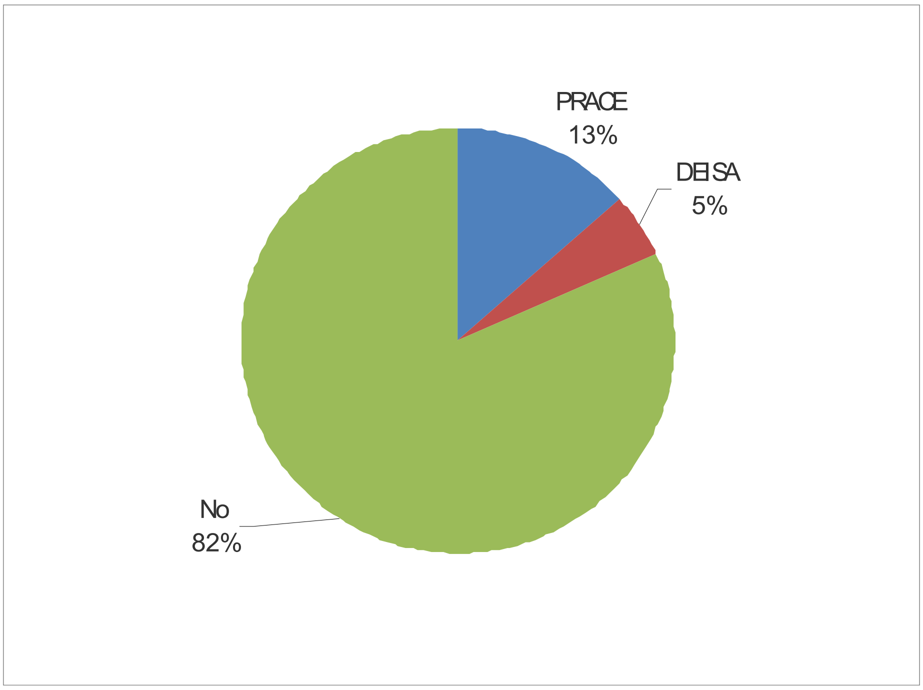

Figure 53: Summary of responses to Question 44.

Question 45: Have you considered applying for resources to PRACE?

Response rate: 24%

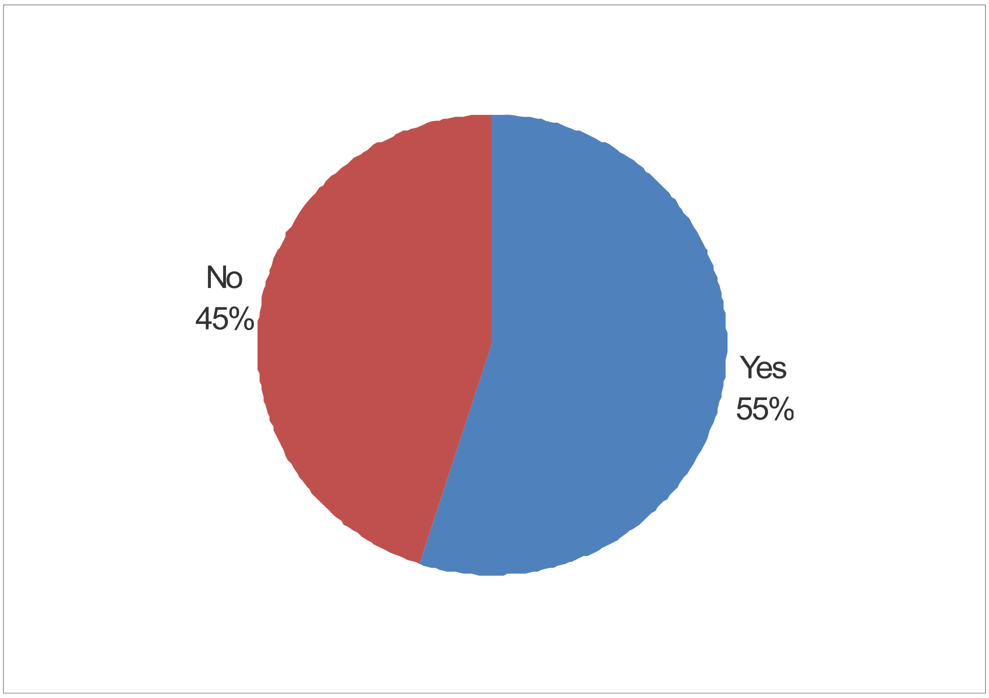

Figure 54: Summary of responses to Question 45.

Question 46: Relating to question 45 above, if not, why not (please choose any that apply)?

Response rate: 44%

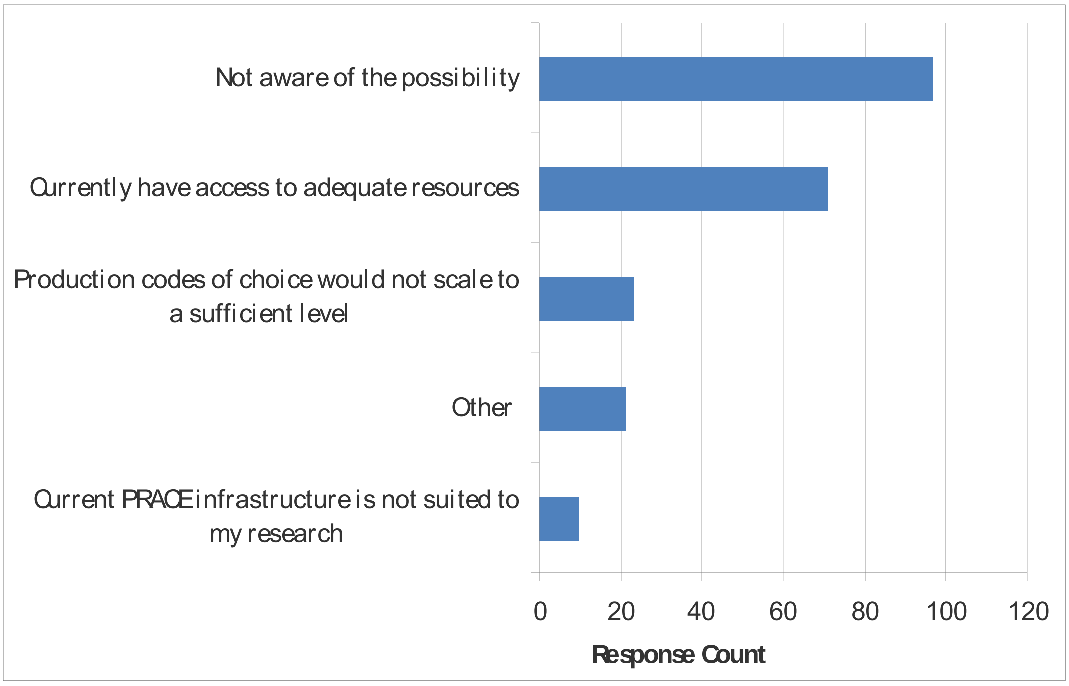

Figure 55: Summary of responses to Question 46.

Comment: The responses to this question clearly show a need to raise awareness of HPC access through PRACE.

Question 47: Will the future availability of different system architectures through PRACE make it more likely that you will apply to use PRACE resources?

Response rate: 44%

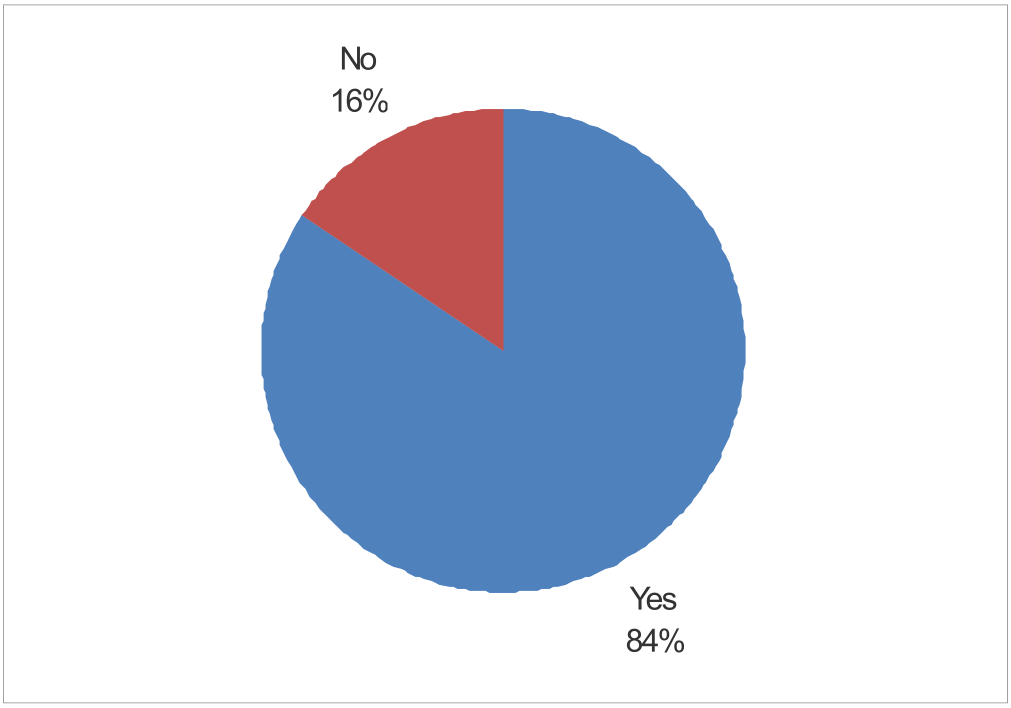

Figure 56: Summary of responses to Question 47.

Question 48: Would the availability of systems that are smaller than the aforementioned Petaflop scale Tier 0 systems, that could act as an intermediary step from your existing production facilities make it more likely that you would apply for Tier 0 access?

Response rate: 63%

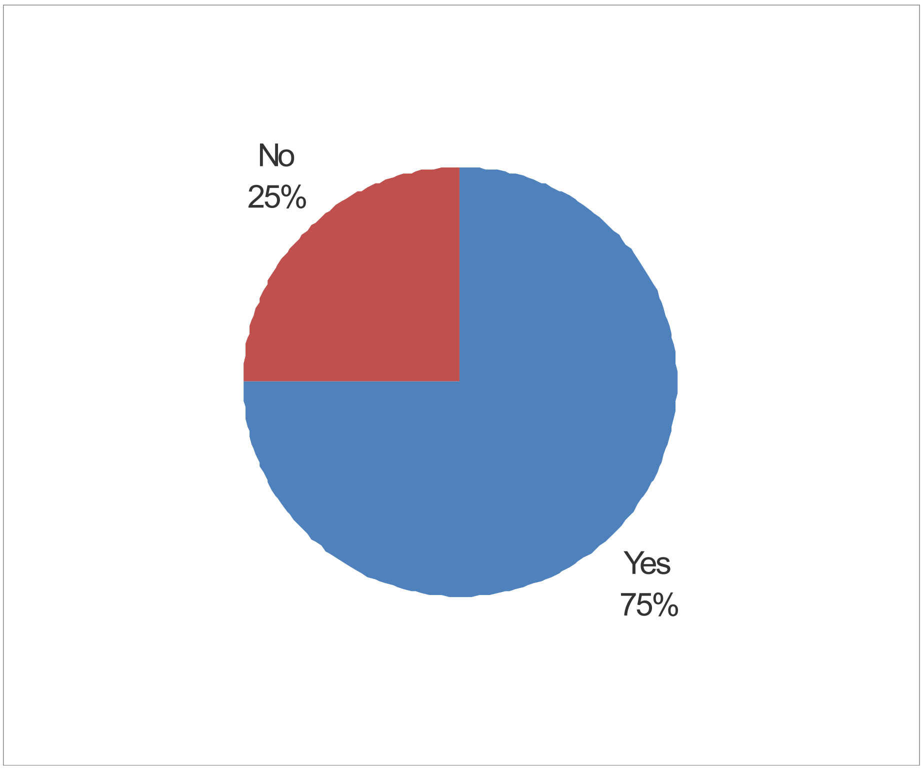

Figure 57: Summary of responses to Question 48.

Question 49: Relating to Question 48 above, would you prefer to use such intermediate systems for: mainly production with some development and scaling, or mainly development and scaling with some production?

Response rate: 53%

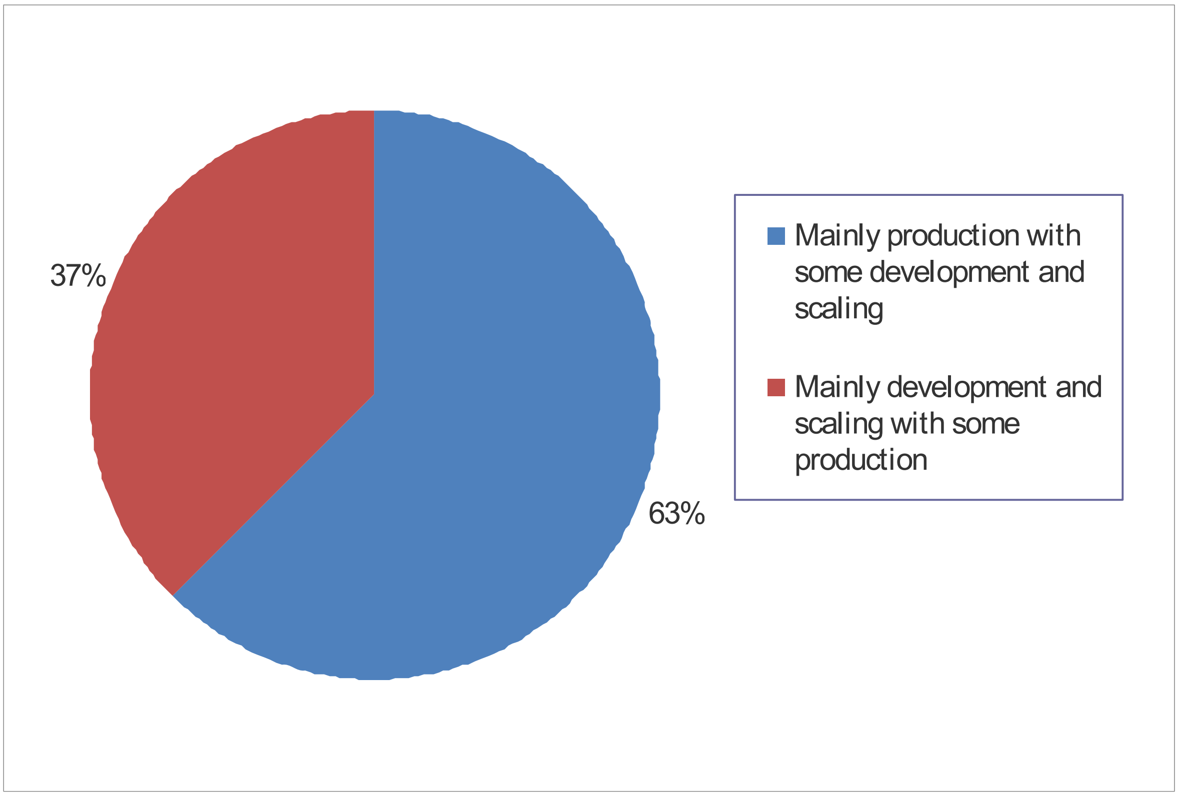

Figure 58: Summary of responses to Question 49.

Question 50: Please let us know your expectations and comments on the PRACE infrastructure and services it offers/will be offering.

Response rate: 16%

Comment: This question was answered with a wide variety of responses, which are difficult to summarise here, as no significant recurring themes were detectable.

4 Conclusions

We have carried out surveys of current PRACE partners' HPC systems, the applications running on them, and of current/potential users of the PRACE infrastructure. These surveys have received a good response rate, and have generated a significant amount of data. This deliverable provides a top-level summary of the survey results, with comparisons to previous survey data where appropriate. The data from these surveys will be made available to other tasks/workpackages in PRACE for further analysis, as required. The following sections summarise the principal findings from the surveys.

4.1 Systems and applications survey

Twenty-eight systems were included in the systems survey, representing just over 3.0 PFlop/s of peak computing power. Compared to the 2008 PRACE-PP survey, the compute power available across the PRACE partners has increased by a factor of 3.7, with almost all this increase being a result of increased numbers of cores, rather than increased power per core. Utilization of systems has increased from 71% in 2008 to 77% in the current survey, but even so, almost a quarter of CPU cycles go unused. It may be worth considering how to better use these wasted cycles, for example by using "bottom-feeding" low-priority jobs.

As the number of cores in systems has increased, so too has the number of cores used by jobs on these systems, but at a slower rate, so that the fraction of the machine used by the average job has decreased since 2008. This indicates than more work is need on the scalability of applications in order to keep pace with the trends in hardware, and to fully transfer the benefits of Moore's Law to application performance. In terms of number of users, the PRACE partner systems are supporting about 50% more users per system than in 2008: a significant increase.

Looking at use by different scientific areas, changes since 2008 are modest. Particle Physics and Computational Fluid Dynamics have higher usage in the current survey, while the proportion consumed by Condensed Matter Physics and Computational Chemistry has decreased.

The survey included an analysis of installed software on the systems. While the popular scientific and I/O libraries are present on many systems, it is clear that they are not universally available. For the PRACE infrastructure of the future, a more consistent and homogeneous environment across systems would be of benefit to users, to prevent users from having to install these libraries themselves, which is time consuming, and may result in sub-optimal installations. The same comments can also be applied to debuggers and especially performance analysis tools, where there is little consistency across systems. The usage of these tools is notoriously low, partly because of the effort required to learn to use them, but this might be addressed at least in part by providing the same (sub)set of tools across all the systems in the PRACE infrastructure.

For each system, we requested an application survey return for all applications which consume more than 5% of the CPU cycles on that system. This resulted in 93 survey returns, representing 57 different applications, and 31% of the total usage of the systems. The applications which consume most compute power are of two types: applications running at large scale on one the largest systems (principally Jugene) and applications which are run on a significant number of different systems. The most widely used codes (in terms of number of systems on which they consume more the 5% of the cycles) are largely the same ones as in the 2008 survey. As a whole, the surveyed applications represent most of the main scientific areas, but Bioinformatics and Plasma Physics are not represented in the survey set.

Compared to the 2008 survey, there has been an increase in the proportion of applications using C or C++, such that the balance is now approximately half C/C++ and half FORTRAN. MPI remains by far the most popular parallelisation technique, though there has been a modest increase in the number of codes using mixed-mode MPI/OpenMP. There were no applications reported as using any of the PGAS family of APIs. FFTW and dense linear algebra (BLAS/LAPACK/ScaLAPACK) are the most heavily used libraries.

4.2 User survey

411 users responded to this survey. The geographical spread is somewhat biased to a small number of countries, and as such the sample may not be very representative of the PRACE partner user population as a whole.

Over 50% users have their own application codes, and consider themselves as developers rather than end users. The most commonly used shared codes are mainly from the areas of Computational Chemistry and Molecular Dynamics.

There is a general desire from users to increase the scalability of their applications, but their ambitions are relatively modest: on average an increase of less than an order of magnitude appears to be sufficient and both strong and weak scaling are important. Over one-third of users do not fully understand the scalability issues of their applications and about 15% expressed a desire for assistance from PRACE to solve this problem. Over one-third of users require more than 2GB of memory per core for their application: this is likely to become a significant problem on future systems. Almost 50% of users would be prepared to use a different application if it meant more scalability.

In terms of future architectures, increased flop rate of CPUs is perceived as the most important architectural feature. Requirements for disk space and for the amount of data to be transferred off the system vary widely, over more than three orders of magnitude.

Echoing the application survey results, users are slightly more likely to use C or C++ than FORTRAN as their main development language, and MPI and OpenMP are by far the most widely used parallelization methods. In general, the use of Grid middleware and workflow systems is rather low.

More than one in six users had attended a PRACE or DEISA training event. Just over 50% of users had considered applying for PRACE resources, but of those that had not, over half were unaware of the possibility, suggesting that further efforts are require to raise awareness. However, a large majority of users thought that the potential future diversity of architectures would make applying for PRACE resources more attractive, and that smaller, stepping-stone resources (i.e. Tier-1) would be a helpful route to Tier-0 usage.

5 Annex A: HP-SEE Project User Survey

5.1 Introduction

The material provided below is the result of the HP-SEE project's user survey that was conducted among the partners of the HP-SEE project in November 2010. It is a summary of the project's deliverable D4.1 "Target Application Analysis" to be provided to other related projects and communities.

The HP-SEE project (http://www.hp-see.eu/) "High-Performance Computing Infrastructure for South East Europe's Research Communities" brings together the National HPC infrastructures in the region of South Eastern Europe and the regional Virtual Research Communities of Computational Physics, Computational Chemistry and Life Sciences, aiming at enabling those user communities to get access to HPC resources for their scientific work.

The following sections provide a summary of the requirements of the target regional applications to be deployed during the project. The relevant survey that revealed the data presented below was targeted to the scientific groups that participate in the HP-SEE project and does not illustrate an exhaustive list of possible HPC applications in the region.

The countries that participate in the project are: Greece, Bulgaria, Romania, Turkey, Hungary, Serbia, Albania, Bosnia-Herzegovina, FYR of Macedonia, Montenegro, Moldova (Republic of), Armenia, Georgia, Azerbaijan. Of those Greece, Bulgaria, Turkey and Serbia are currently members of PRACE and have also responded to the WP7 questionnaire while the others are not.

5.2 The HP-SEE VRC and Applications

HP-SEE supports and strengthens a number of strategic Virtual Research Communities, which will bring together users across the region within a common cooperative research space, enabling them to share HPC facilities, software, tools, data and results of their work. Thus, the project will directly contribute to the co-ordination of high-quality research and ease the access to and enhance the usability of the available infrastructure. The core international scientific fields identified as self-standing Virtual Research Communities are Computational Physics, Computational Chemistry and Life Science Virtual Community.

5.3 Computational Physics community

Computational physics is nowadays the main beneficiary of the scientific HPC, large-scale numerical computations being necessary whenever the complexity of the physical systems investigated does not allow the derivation of an analytical solution.

The main objective of the Computational Physics VRC is to join together the various physics research teams from the SEE area and to provide them access to a powerful HPC infrastructure and tools which will make possible their participation in multidisciplinary and international collaborations.

For this purpose, software developers from 6 countries will contribute with 8 applications in the fields of High Energy and Particle Physics, Plasma Physics, Physics of Condensed Matter, Atomic Physics, and Computational Fluid Dynamics.

The application range extends from nanoelectronics, micro-devices optimization and the modeling of robotic devices for biomedicine, to improved means for feature detection in satellite images, which leads to better mapping, localization and search services.

5.4 Computational Chemistry virtual research community

Computational chemistry and material science is one of the highlighted research areas in computational science and a typical heavy user of HPC resources. The computational technologies are an indispensable tool for investigations in domains like quantum molecular dynamics, molecular modelling, nano-technology and design of new materials. Considering the size of the problems to be studied, the required calculations are extremely computationally intensive. Thus HPC would greatly facilitate the proposed work allowing the researchers to deal not only with “pilot” or model systems but to work on big and complicated real systems, which are physically and technologically more significant and challenging. These studies will extend understanding of some fundamental science issues and are of practical importance for pharmaceutical industry, nanotechnology, biomedicine, and many others.

Initially Computational Chemistry VRC supports 7 applications with main developers in 6 SEE countries, collaborating with scientists from more than 20 advanced research centers in Europe.

5.5 Life Sciences community

Life Sciences depend heavily on the use of HPC for both data mining and data integration as well as for the simulation of biological systems. HPC technologies are essential for research areas such genome analysis, expression profiling, -omics analysis and biological simulations, whereby a vast amount of experimental data needs to be analyzed and synthesized into reasonable hypothesis. Thus HPC would greatly facilitate the various applications described in this project, enabling the respective research teams to study questions that have thus far been intractable due to their high computational complexity. The use of HPC in the Life Sciences applications will help further our understanding of basic problems in the fields of DNA sequence analysis, comparative genomics, and brain modeling among others and can be of great importance for the health sector.

The Life Sciences VRC supports 7 applications with main developers in 5 SEE countries (Greece, Hungary, Montenegro, Armenia, Georgia) working in the areas of computational biology, computational biophysics, DNA sequence analysis and computational genomics. The various projects involve collaborations with numerous scientists both in Europe and the U.S. and will foster the development of new collaborations among the participant SEE countries.

5.6 Supported Applications

A total number of 26 candidate applications were suggested, via the replies to the users' survey, 22 of which were planned to be supported by the project. The Computational Chemistry VRC and the Life Sciences VRC have contributed with seven applications each, while 12 applications were grouped in the Computational Physics VRC. It is to be noted that the number of applications should not be interpreted as a measure of the VRC size, as the complexity of the applications and the number of their beneficiaries can vary considerably.

The tables below capture the list of applications on each VRC, together with the main developer institution, the development stage at the time of the survey, and the estimated time when the application will be ready for production.

The contributions to the Computational Chemistry VRC are provided by universities and research institutes from six Balkan countries and, with one notable exception, their production phase is planned after the first year of the project.

Acronym	Application name	Main developer
CFDOF	CFD Analysis of Combustion	Faculty of Mech. Engineering, University of Banja Luka (UoBL), Bosnia - Herzegovina
CompChem	Quantum Mechanical, Molecular Mechanics, and Molecular Dynamics computation in chemistry	Univeristy of Belgrade, Faculty of Chemistry
FMD-PA	Design of fullerene and metal-diothiolene-based materials for photonic applications	Computational Chemistry Group of NHRF. Greece
HC-MD-QM-CS	Hybrid Classical/Quantum Molecular Dynamics – Quantum Mechanical Computer Simulation of Condensed Phases	UKIM, Institute of Chemistry, Faculty of Natural Science and Mathematics, FYROM
ISyMAB	Integrated System for Modeling and data Analysis of complex Biomolecules	IFIN-HH/DPETI, Romania
MDCisplatin	Molecular Design of Platinum Group Metal Complexes as Potential Non-classical Cisplatin Analogues	Acad. Roumen Tsanev Institute of Molecular Biology, Bulgarian Academy of Science
PCACIC	Principal component analysis of the conformational interconversions in large-ring cyclodextrins	IOCCP-BAS, Bulgaria

Table 6: Computational Chemistry Applications

The main applications developers of the Computational Physics VRC are from seven Balkan countries and, according to the estimations, three of their applications can be ready for production during the first year of the project.

Acronym	Application name	Main developer
AMR_PAR	Parallel algorithm and program for the solving of continuum mechanics equations using Adaptive Mesh Refinement	Inst. of Mathematics and Computer Science (IMI ASM)/ Laboratory of Mathematical Modeling, Moldova
EagleEye	Feature Extraction from Satellite Images Using a Hybrid Computing Architecture	University Politehnica of Bucharest (UPB) / Computer Science and Engineering, Romania
FAMAD	Fractal Algorithms for MAss Distribution	Institute of Space Sciences (ISS), Romania
FuzzyCmeans	Parallel Fuzzy C Means for classification/Feature detection category	West University of Timisoara (UVT) /Computer Science Department, Romania
GENETATOMIC	Genetic algorithms in atomic collisions	UKIM, FYROM
GIM	Geophysical Inversion Modeling	Polytechnic University of Tirana
HAG	High energy physics Algorithms on GPU	Institite of Space Sciences (ISS), Romania
HMLQCD	Hadron Masses from Lattice QCD	Univeristy of Tirana
NUQG	Numerical study of ultra-cold quantum gases	SCL, Institute of Physics Belgrade
SET	Simulation of electron transport	IICT-BAS, Bulgaria
SFHG	Self Avoiding Hamiltonian Walk on Gaskets	Faculty of Mechanical Engineering, Dept. of Thermomechanics, University of Banja Luka, Bosnia and Herzegovina
SIMPLE-TS 2D	Finite Volume Method for calculation of 2D gas-microflows using standard MPI	Institute of Mechanics – BAS / Dept. of “Complex and multiphase Flows”, Bulgaria

Table 7: Computational Physics Applications

The developer community with the widest geographical distribution is that of the Life Sciences, which is hosted in five countries and foresees the provision of 7 applications, from which one during the first year of HP-SEE.

Acronym	Application name	Main developer
CMSLTM	Computational Models of Short and Long Term Memory	IMBB-FORTH, Greece
DeepAligner	Deep sequencing for short fragment alignment	Obuda University (OU), John von Neumann Faculty of Informatics, Biotech Group, Hungary
DiseaseGene	In-silico Disease Gene Mapper	Obuda University (OU), John von Neumann Faculty of Informatics, Biotech Group, Hungary
DNAMA	DNA Multicore Analysis	School of Computer & Communic. Sciences, Laboratory for Computational Biology and Bioinformatics (LCBB), Montenegro
MDSCS	Molecular Dynamics Study of Complex systems	IIAP NAS, Armenia
miRs	Searching for novel miRNA genes and their targets	IMBB/FORTH, Greece
MSBP	Modeling of some biochemical processes with the purpose of realization of their thin and purposeful synthesis	Tbilisi State University Department of Natural Science, Georgia

Table 8: Life Sciences Applications

5.7 Profile of HP-SEE project user communities

5.7.1 *Basic aspects*

Figure 59 shows the distribution of the 26 supported applications according to the contributing countries.

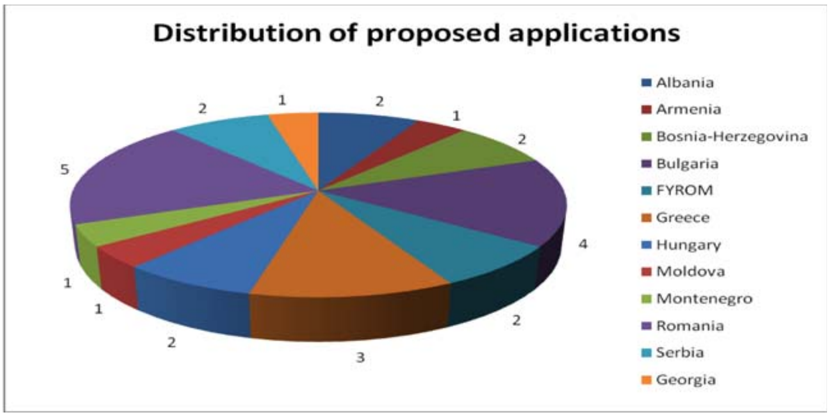

Figure 59 : Distribution of the proposed applications according to countries.

According to the statistics below, the primary programming language is C/C++ for 13 applications, being followed by Fortran (10 cases, including one Intel Fortran), Java (2 applications), and NMODL - a high-level language used in the NEURON simulation environment.

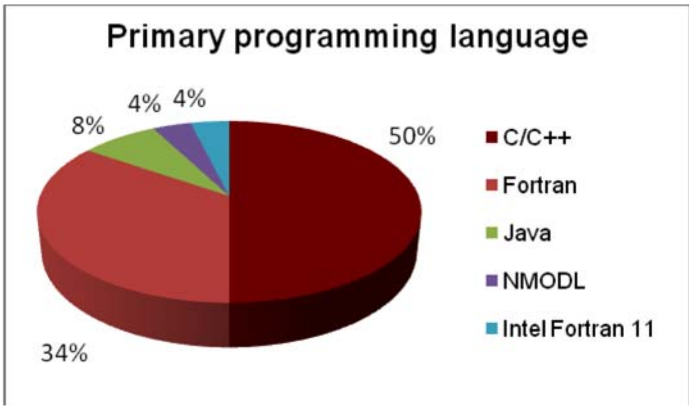

Figure 60: Distribution of applications according to the primary programming languages used.

The most used compilers are the free GCC, which is preferred for 20 applications, and Intel, used for 13 applications. Minor shares are taken by Matlab and PGI (with 3 applications each), and javac (1 application).

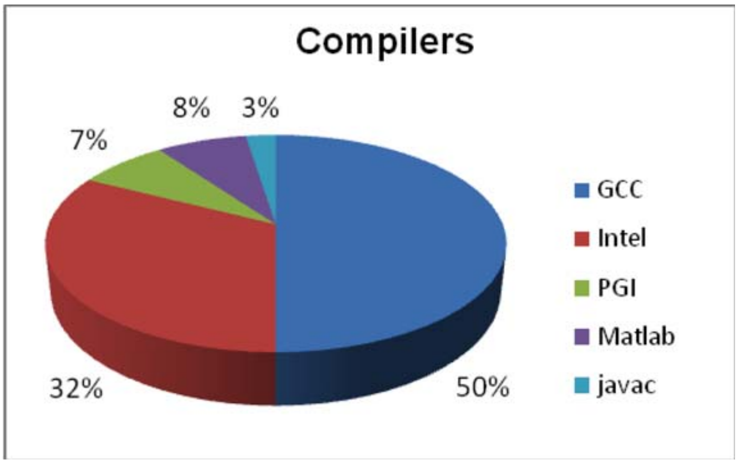

Figure 61: Percentage distribution of the compilers

A fundamental question which was addressed to the developers concerned the paradigm of the parallel computing which is used by programs. 43% of the responses have designated MPI (clustered multiprocessing), followed by 21% using Shared Memory Programming (OpenMP), 19% using CUDA or OpenCL, and 17% which use the clusters for multiple serial jobs.

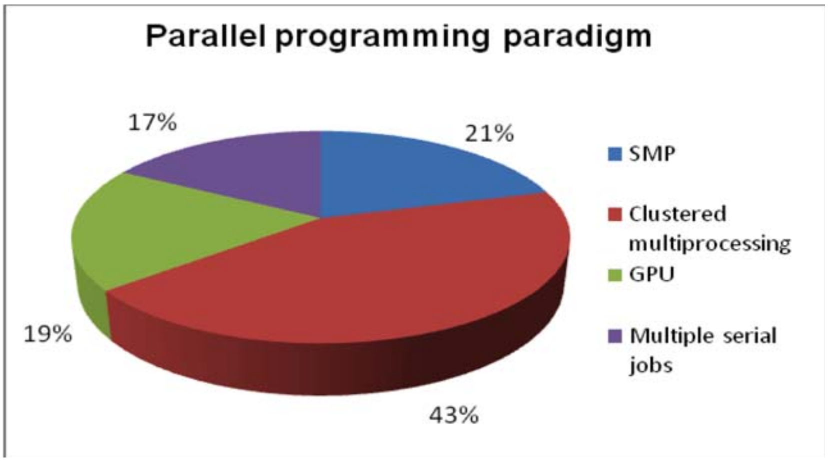

Figure 62: Percentage distribution of the used types of parallel programming

5.7.2 *Libraries*

Note that in this section, the term "libraries" includes codes which are used as third-party applications.

The ab-initio quantum chemistry package GAMESS leads the group of the scientific libraries, being necessary to four Computational Chemistry applications and one Plasma Physics program.

This is closely followed by the classical and ab-initio molecular dynamics programs NAMD and CPMD, respectively, which are used by groups from the Computational Chemistry and Life Sciences VRCs, and the GotoBLAS library for linear algebra.

The group of the libraries which are used by two HP-SEE applications is composed of: OpenFOAM package for fluid dynamics; the proprietary code Gaussian - for electronic structure modeling; BioPerl - for bioinformatics, genomics and life sciences in general; the library of random number generators SPRNG; NWChem - for computational chemistry; the ROOT framework for large scale data analysis.

Finally, there is a group of 10 libraries, each of which is used by only one application; these are listed in Figure 63 below.

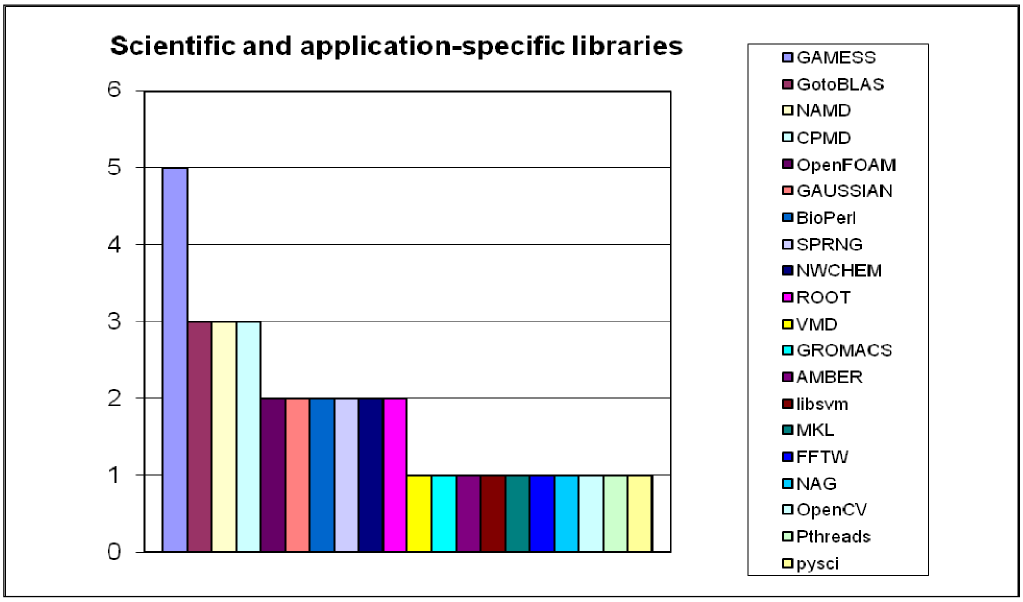

Figure 63: The use of the scientific and application-specific libraries

The set of libraries required for input/output operations, for communication, or for general purposes is represented in Figure 64 and is dominated by the free implementations of MPI MPICH-2 and MPICH-1, followed by OpenMP for SMP; their relative weight is correlated to the distribution of parallel computing paradigms from Figure 62. Next come the linear algebra libraries (BLAS, ATLAS), the discrete fast Fourier transform library FFTW, the mathematical MKL library from Intel, and the linear algebra library ScaLapack. ROOT library is used by the two CPVRC applications proposed by ISS.

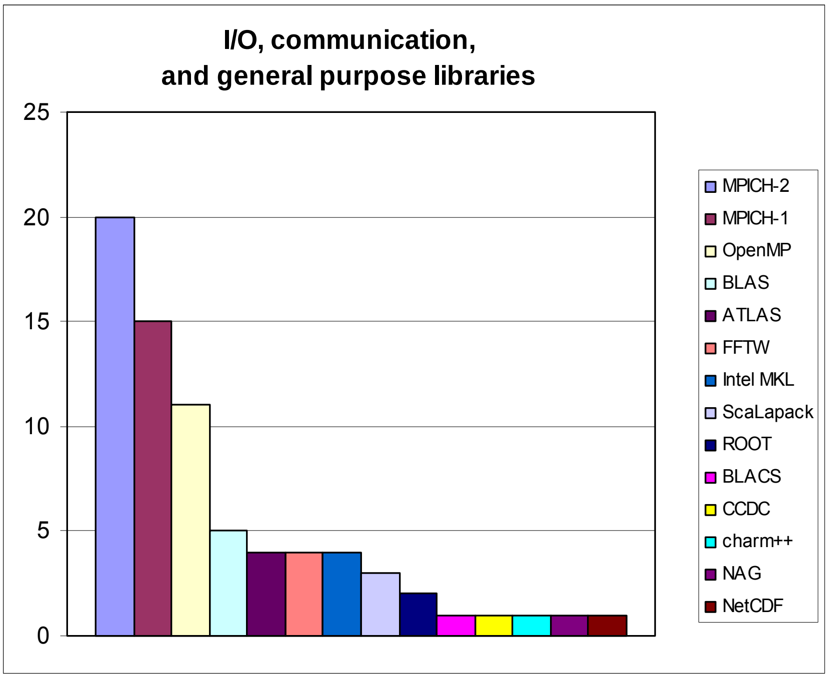

Figure 64: The usage of I/O, communication, and general purpose libraries

5.7.3 *System Requirements*

Before analyzing the system requirements of the applications, it is useful to know on what initial systems the development was performed at the time of the survey. The answers revealed that two thirds of the applications were hosted on home clusters, while 5 applications still used workstations and one the Grid.

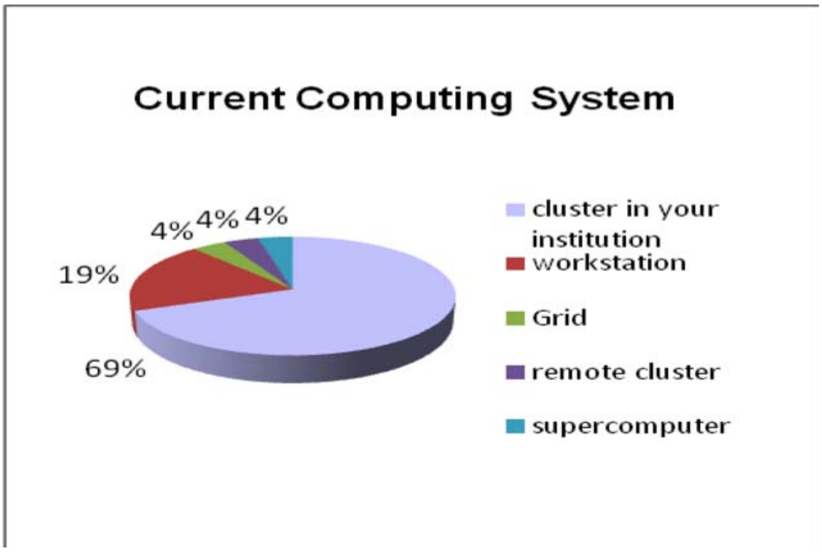

Figure 65: Distribution of the applications according to the type of the currently hosting system

With respect to the CPU requirements, we remark that the 64 bit and high-end CPUs are preferred. It is to be noted that there are plans to port one application from vector processing to GPU.

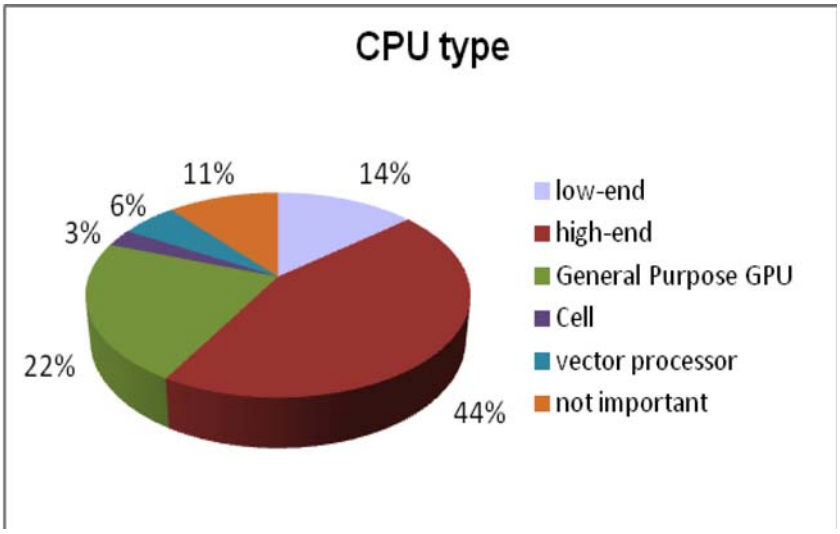

Figure 66: Distribution of the applications according to the CPU architecture and type

The first basic resource requirement for a parallel code regards the number of logical processors (cores). Figure 67 and Figure 68 present the estimated number of cores that is necessary in the initial stage of the applications and in the production regime. While in their initial stage half (13) of the applications need at most 64 cores, in the production stage 14 of them require 128-512 cores to run. Also, in production CompChem and NUQG will require up to 4096 cores, while HMLQCD and SET will both need more than 4096 (8000 cores for SET).

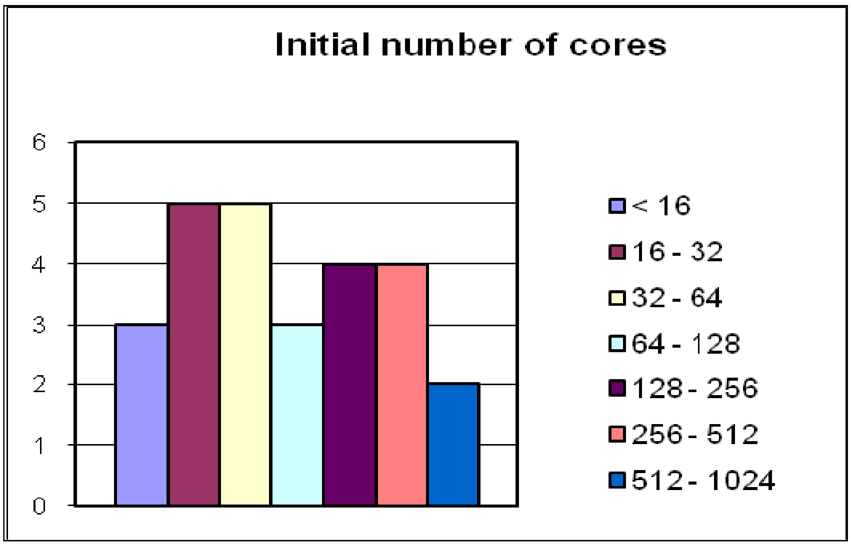

Figure 67: Distribution of the applications according to the number of cores currently in use or initially desired

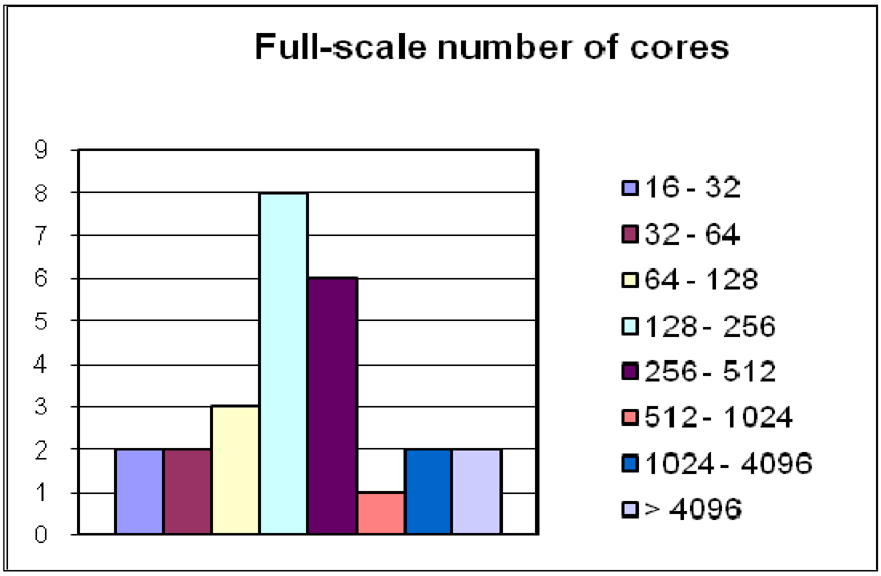

Figure 68: Distribution of the applications according to the number desired number of cores at full scale

5.7.4 *Storage requirements*

The third basic parameter is the storage capacity required by the application. In what follows we present the statistics of the estimations regarding the space the applications need for temporary storage, installation, and long-term storage.

Most of the applications will initially need for one run 1-2 GB; the maximum initial value is 1TB, which is necessary for one application. Most of the applications will require in the production regime 10 GB or less, but the upper bound of 1 TB will not change.

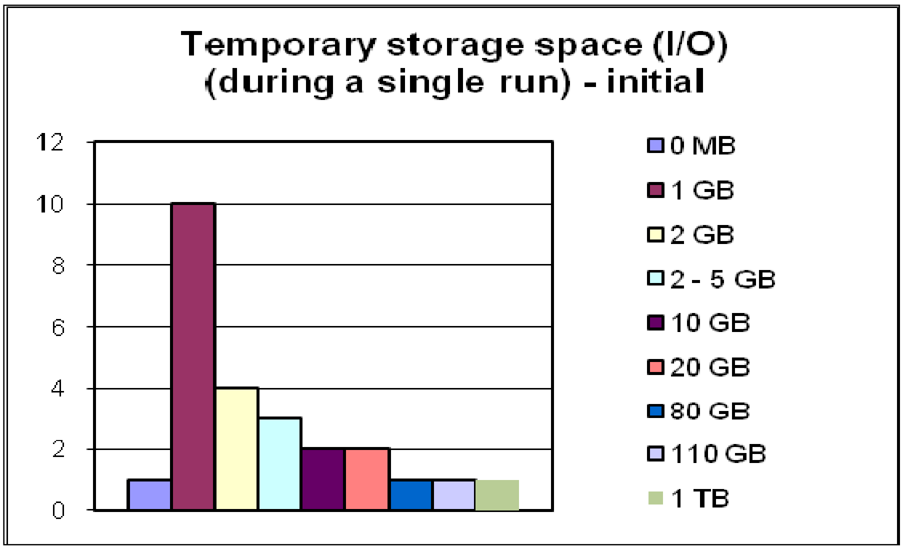

Figure 69: Distribution of the applications according to the temporary storage space currently in use or initially desired

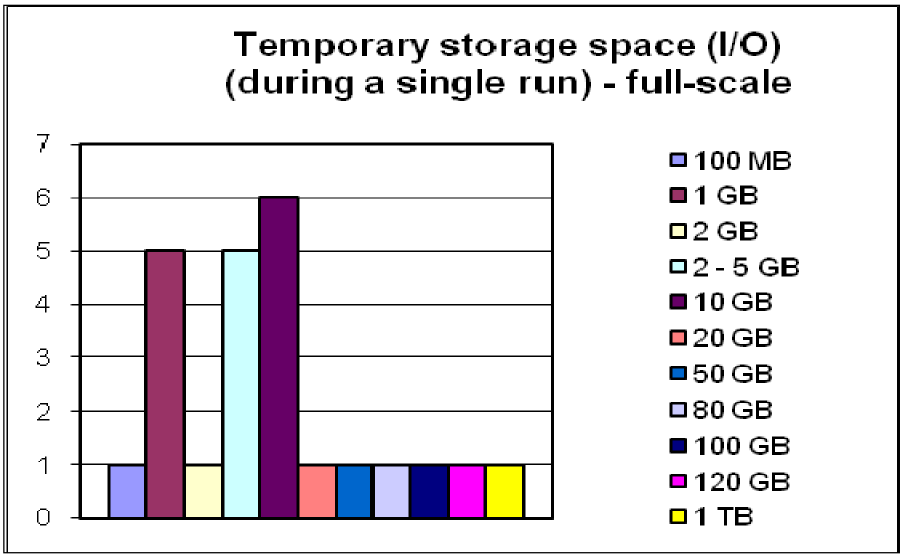

Figure 70: Distribution of the applications according to the temporary desired storage space at full scale

20 applications initially need less than 500 GB for long-term storage, and 16 of them will kep the same range of values at full scale. The extrema are reached by one application, which needs initially 0 MB and requires 25 TB in the production regime.

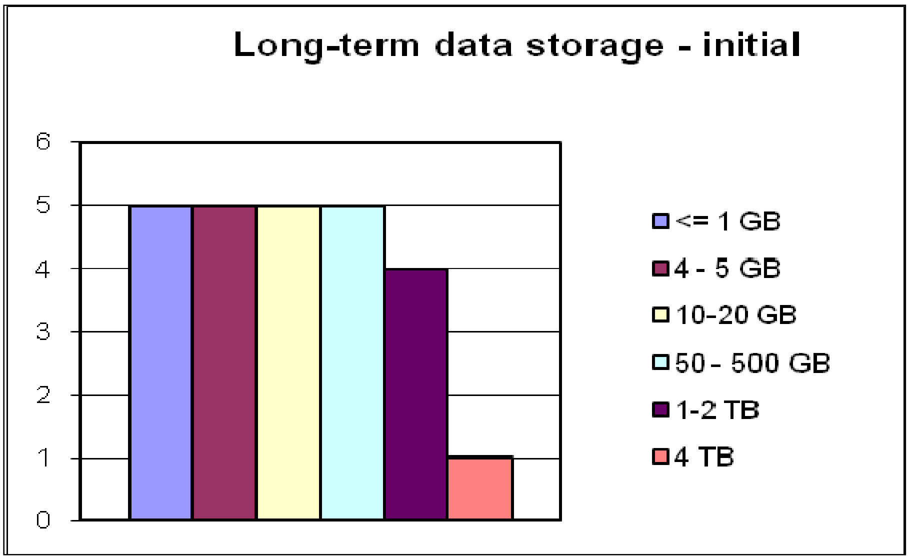

Figure 71: Distribution of the applications according to the long-term data storage currently in use or initially desired

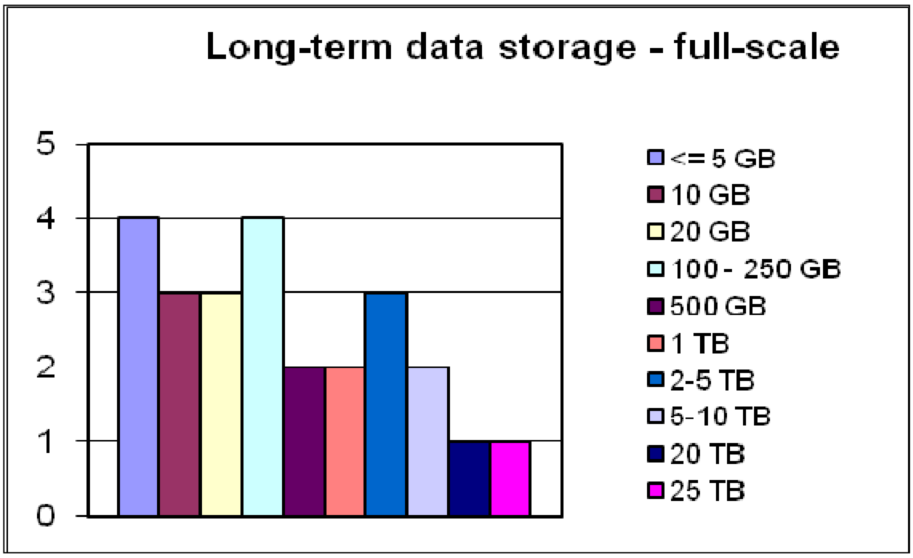

Figure 72: Distribution of the applications according to the long-term data storage at full scale

Another parameter which is important for both users and resource providers is the minimum duration of the long-term storage. Statistics shows that, when passing from the initial to the production stage the mean storage time increases; if initially 9 applications require storage for more than 3 months, their number grows to 13 in the full-scale regime. The reasons why some applications need a minimum duration of two years for long-term storage should be elucidated.

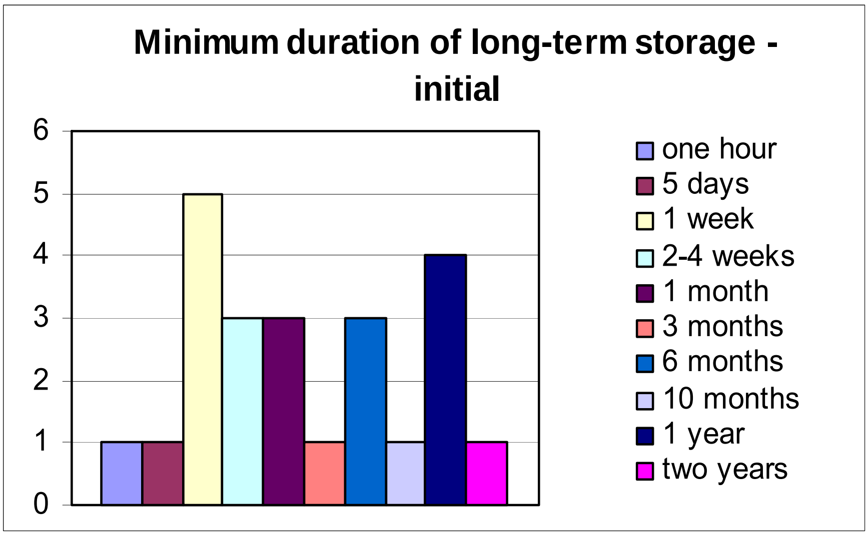

Figure 73: Distribution of the applications according to the required storage duration currently in use or initially desired

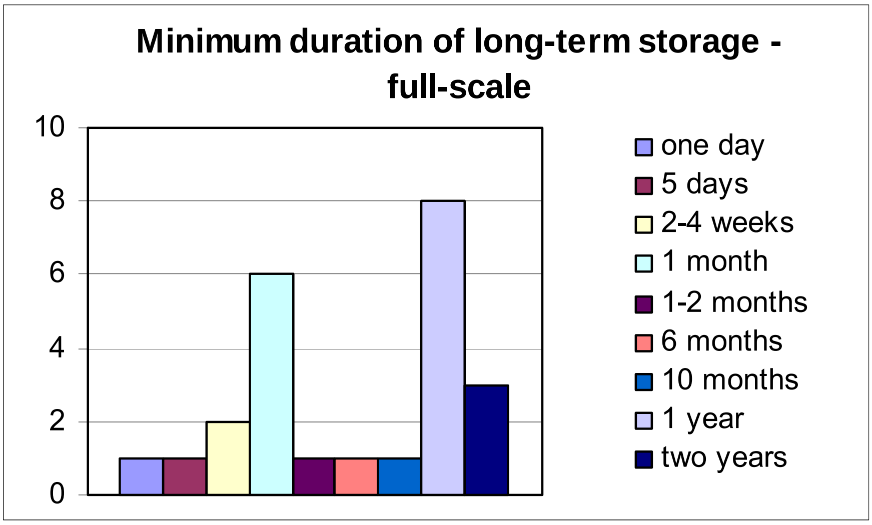

Figure 74: Distribution of the applications according to the required storage duration at full scale

According to the survey, all the applications are Linux-compatible, and many of them with Unix too. 5 applications are compatible with Linux, Unix and MS Windows.

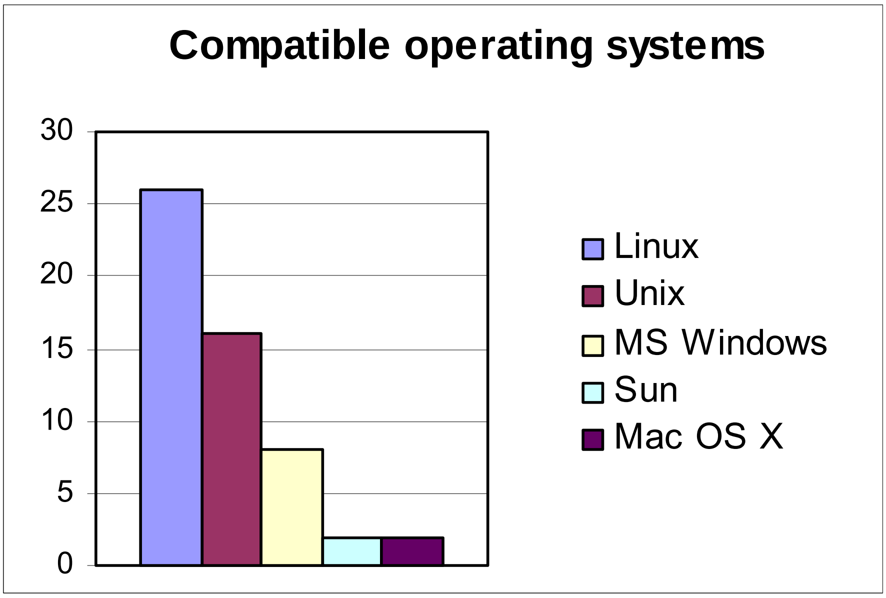

Figure 75: Distribution of the applications according to the compatible operating systems

5.7.5 *Conclusions*

This Annex is a summary of HP-SEE Deliverable D4.1 “Target Applications Analysis” and presents an overview of the HP-SEE target applications, based on the information gathered during the first quarter of the project by means of the online applications questionnaire. The collected data concern the applications requirements, the user communities, and the training needs. In this summary we present only a summary of the applications requirements.

The analysis of these data revealed a detailed picture of the three virtual research communities that support the applications, within which many research groups already reached an advanced level of collaboration and new cooperation is initiated by the partners of the HP-SEE project.