



**SEVENTH FRAMEWORK PROGRAMME
Research Infrastructures**

INFRA-2007-2.2.2.1 - Preparatory phase for 'Computer and Data Treatment' research infrastructures in the 2006 ESFRI Roadmap



PRACE

Partnership for Advanced Computing in Europe

Grant Agreement Number: RI-211528

D2.6.1

Operational model analysis and initial specifications

Final

Version: 1.3
Author(s): S. Requena, GENCI
Miroslaw Kupczyk, PSNC
Date: 30.12.2008

Project and Deliverable Information Sheet

PRACE Project	Project Ref. №: RI-211528	
	Project Title: Partnership for Advanced Computing in Europe	
	Project Web Site: http://www.prace-project.eu	
	Deliverable ID: < D2.6.1 >	
	Deliverable Nature: Report	
	Deliverable Level: PU	Contractual Date of Delivery: 31/12/2008
		Actual Date of Delivery: 31/12/2008
EC Project Officer: Maria Ramalho-Natario		

* - 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: Operational model analysis and initial specifications	
	ID: D2.6.1	
	Version: <1.3 >	Status: Final
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2003	
	File(s): D2.6.1.doc	
Authorship	Written by:	Mirek Kupczyk, PSNC Stephane Requena, GENCI
	Contributors:	A. Lichnewsky, GENCI Sergi Girona, Eugene Griffith, BSC Emma Jones, EPSRC Francois Robin, GENCI/CEA Jacko Koster, UNINETT Sigma Stefan Wesner, HLRS Radek Januszewsk, PSNC Agnieszka Kwiecien, PSNC
	Reviewed by:	Thomas Eickermann, FZJ Peter Michielse, NCF
	Approved by:	Technical Board

Document Status Sheet

Version	Date	Status	Comments
0.1	30/06/2008	Draft Outline	After discussion in BCN, WP2.6 Meeting
0.2	14/07/2008	Draft Outline+some Preliminary TOC,	Alain, Stephane
0.3	5/10/2008	Introductory, methodology	Stephane, Mirosław
0.4	24/11/2008	Draft	Final responses collected and processed - Mirosław, Agnieszka, Radek, Stephane
0.5	26/11/2008	Draft	Processing the responses cont. Mirosław, Agnieszka, Radek, Stephane
0.6	03/12/2008	Draft	Jacko - ecosystem, Francois – management Agnieszka – improved BOffice, Radek – Security&HPC
0.62	05/12/2008	Draft	Agnieszka – update Back Office&Training Mirek – merging Polices+BOffice, HPC-still draft
0.64	05/12/2008	Draft	Merge between 0.62 and Stephane modifications
0.65	07/12/2008	Draft	Stephane – Survey results
0.7	09/12/2008	Draft	Integration of Radek, Mirek and Jacko modifications
1.0	09/12/2008	Version for review	
1.1	15/12/2008	Draft	Stephane incorporates the remarks
1.2	18/12/2008	Draft	Mirosław incorporates the remarks from reviewer.
1.3	23/12/2008	Final	Stephane: Final integration of all the remarks

Document Keywords and Abstract

Keywords:	PRACE, HPC, Research Infrastructure
Abstract:	<p>PRACE will provide the research community with a single, unique permanent infrastructure, thus permitting and encouraging users to invest in the development of simulation codes based on computational sciences.</p> <p>A consistent operational model is thus required for PRACE to achieve this goal, in the face of evolving technologies and architectures, across multiple interconnected distributed sites. It must also allow for building a coherent ecosystem ensuring that PRACE tier-0 systems are easy to use for users familiar with the other 2 tiers of the European HPC environment.</p> <p>This document presents an analysis of the operational model for PRACE and initial specifications.</p>

Copyright notices

© 2008 PRACE Consortium Partners. All rights reserved. This document is a project document of the PRACE project. All contents are reserved by default and may not be disclosed to third parties without the written consent of the PRACE partners, except as mandated by the European Commission contract RI-211528 for reviewing and dissemination purposes.

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

Table of Contents

Project and Deliverable Information Sheet	i
Document Control Sheet.....	i
Document Status Sheet	ii
Document Keywords and Abstract.....	iii
Table of Contents	iv
References and Applicable Documents	v
List of Acronyms and Abbreviations.....	vi
Executive Summary	1
1 Introduction	2
2 Organisation of this document	4
2.1 Assumptions.....	4
3 Front Office.....	6
3.1 HPC	9
3.2 Related services.....	12
3.2.1 Data.....	12
3.2.2 Archive.....	13
3.2.3 Graphics/Visualisation.....	13
3.2.4 Pre- and Post- processing	14
3.2.5 Interactive service	14
3.3 Initial specifications for Front Office services	15
4 Ecosystem integration	16
5 Training, support and dissemination.....	22
5.1 User support.....	22
5.2 User training	24
5.3 Personnel training	26
5.4 Dissemination and outreach	26
5.5 Vendor Documentation & training with regard to vendor tools.....	27
5.6 Initial specifications for the Training, Support and Dissemination services.....	28
6 Security.....	29
6.1 Initial specifications for the Security services	33
7 Back Office.....	34
7.1 Handling of exceptional resource requests.....	34
7.2 Main processes.....	37
7.3 Main parts of organisation.....	38
7.4 System administration.....	39
7.5 Critical delivery time.....	40
7.6 Initial specifications for the Back office services	40
8 Management and decision support.....	41
8.1 Management and decision support	41
8.2 Planning and procurement management	43
8.3 Monitoring and reporting.....	44
8.4 Initial specifications for Management and decision support	45

9	Conclusion and next steps.....	46
10	Annex.....	47
10.1	Description of the survey	47
10.2	Survey answers	47
10.3	Survey details and results	48
	Introduction	48
	HPC & Related Services	48
	Ecosystem integration	56
	User Support.....	58
	Training	60
	Dissemination.....	63
	Data security.....	64
	Site Security	66
	Operation.....	67
	Back office	70
	Management and decision support.....	71

References and Applicable Documents

- [1] <http://www.prace-project.eu>
- [2] PRACE Project, D2.4.1 “Initial Report on the Peer Review Process”.
- [3] PRACE Project, D3.1.1 “Final plan for the use and dissemination of foreground”.
- [4] PRACE Project, D3.3.1 “Survey of HPC education and training needs”.
- [5] PRACE Project, D3.3.2 “ PRACE Summer School”.
- [6] PRACE Project, D3.3.3 “PRACE Winter School”.
- [7] PRACE Project, D2.1.1 “Report on options for a legal entity”.
- [8] PRACE Project, D2.5.1 “Analysis of the HPC Ecosystem”.
- [9] PRACE Project, D2.3.2 “Usage Model Document ”.
- [10] PRACE Project, D4.1.1 “Requirements analysis for tier-0 system management ”.
- [11] PRACE Project, D4.1.2 “Report on existing tier-0 systems management solutions”.
- [12] PRACE Project, D4.2.1 “Evaluation report of existing solutions for the ecosystem integration”.
- [13] PRACE Project, D7.5.1 “Technical Requirement for the first Petaflop/s systems(s) in 2009/2010”.

List of Acronyms and Abbreviations

AAA	Authorisation, Authentication, Accounting.
COREGRID	European Network of Excellence for Grid and Peer-to-Peer technologies.
DEISA	Distributed European Infrastructure for Supercomputing Applications. EU project by leading national HPC centres.
EGEE	Enabling Grids for E-science; EU Grid project lead by CERN and successfully completed in 2004. Follow-up is EGEE-II.
ESFRI	European Strategy Forum on Research Infrastructures; created roadmap for pan-European Research Infrastructure.
HET	High Performance Computing in Europe Taskforce. Taskforce by representatives from European HPC community to shape the European HPC Research Infrastructure. Produced the scientific case and valuable groundwork for the PACE project.
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing.
HPC-Europa	Consortium of six leading (HPC) infrastructures and five centres of excellence providing transnational access; EU project.
ISO	International Organisation for Standardisation
ITIL	Information Technology Infrastructure Library, a set of concepts and policies for managing information technology (IT) infrastructure, development and operations.
LRMS	Local Resource Management System
MoU	Memorandum of Understanding.
NDA	Non-Disclosure Agreement. Typically signed between vendors and customers working together on products prior to their general availability or announcement.
GÉANT	Collaboration between National Research and Education Networks to build a multi-gigabit pan-European network, managed by DANTE. GÉANT2 is the follow-up as of 2004.
PRACE	Partnership for Advanced Computing in Europe; Project Acronym.
Phosphorus	Pan-European optical network test-bed providing bandwidth on-demand services for Grid applications.
SIRENE	Cooperative association of countries forming a pan-European group and policy framework for strong cooperation towards the initial setup and deployment of interoperable Grid infrastructures.
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.
tier-2	Regional HPC centre
TCO	Total Cost of Ownership. Includes the costs (personnel, power, cooling, maintenance, ...) in addition to the purchase cost of a system.

STRATOS	PRACE advisory group for STRAtegic TechnOlogieS
SLA	Service Level Agreement
SME	Small and Medium Enterprise
UNICORE	Uniform Interface to Computing Resources. Grid software for seamless access to distributed resources.
WP2	PRACE work package in charge of Organisational concept of the Research Infrastructure
WP3	PRACE work package in charge of Dissemination, Outreach and Training.
WP4	PRACE work package in charge of the Distributed System Management.
WP5	PRACE work package in charge of the Deployment of Prototype Systems.
WP6	PRACE work package in charge of the Software Enabling for Petaflop/s Systems.
WP7	PRACE work package in charge of the Petaflop/s Systems for 2009/2010.
WP8	PRACE work package in charge of the Future Petaflop/s computer technologies beyond 2010.

Executive Summary

PRACE will provide the research community with a single, unique permanent infrastructure, thus permitting and encouraging users to invest in the development of simulation codes based on computational sciences.

A consistent operational model is thus required for PRACE to achieve this goal, in the face of evolving technologies and architectures, across multiple interconnected distributed sites. It must also allow for building a coherent ecosystem ensuring that PRACE tier-0 systems are easy to use for users familiar with the other 2 tiers of the European HPC environment.

This document presents through the result of an intensive questionnaire sent during November 2008 to European HPC tier-1 and tier-2 centres and completed by 18 of them an analysis of the model for the technical operation of PRACE. Implementation details of mandatory services like Front Office services (access to HPC resources, data management, visualisation, pre/post processing, archive), Training/User Support/Dissemination, Ecosystem integration (in order to provide a seamless integration of the PRACE tier-0 centres on the existing European HPC ecosystem) Security, Back Office services (administration, peer review, ...) have been gathered in order to have an initial view and specifications of which services will be implemented on PRACE tier-0 systems.

As an initial step, this deliverable will be completed in 2009 following other upcoming deliverables and discussions with technical Work packages by D2.6.2 which will define the “best set” of services available on PRACE tier-0 systems.

1 Introduction

The PRACE project has the overall objective to prepare the creation of a persistent pan-European HPC service. The work of PRACE is organized in eight inter-linked work packages (WP). WP2, “Organisational concept of the Research Infrastructures”, focuses on the preparation of a comprehensive legal and administrative framework for the implementation phase of the research infrastructure. Furthermore WP2 deals with the specification of funding and usage strategies, the establishment of a peer-review process, the establishing of links with the HPC ecosystem, and the development of the operation model.

The Merriam-Webster dictionary defines “operations¹” as “performance of a practical work or of something involving the practical application of principles or processes”. Therefore, the operational model is a description of the services and processes to be put in place to perform PRACE’s tasks.

The Task 2.6 in Work Package 2 has the responsibility to analyse and then specify the “Models for Technical Operation” for PRACE. The result of these activities will be described in two deliverables.

The current one (D2.6.1) is the first of them and will deal with an analysis of existing models for operation of RIs and provide an initial specification. This analysis will be made upon the result of an extensive questionnaire sent to existing European tier-1 (mainly) and some tier-2 HPC Centres.

Our choice for this questionnaire was to target only European HPC Centres (but involved into pan European projects like DEISA, EGEE and HPC Europa or delivering yet an HPC service to distributed users through national grids) instead of enlarging it to others non HPC RI.

For each category of service listed we will give the result of the questionnaire and try to define an initial set of services for the upcoming PRACE RI.

The following D2.6.2, targeted for month 24, will give a complete but not final specification.

It will take as an input the initial description done in 2.6.1 and through work/results of deliverables of PRACE’s technical work packages (especially WP4, WP5, WP7 and WP8) and other upcoming deliverables from WP2 and WP3 will define the “best set” of services delivered by the 3 to 5 HPC infrastructures which will be deployed starting 2010.

The process of designing policy models cannot be stopped at the given time – it is a process, which collects the important points from the HPC ecosystem continuously. D2.5.1 [8] has proposed to set up the collaboration with existing policy makers.

Within the PRACE project, these documents have the role of defining the various services and processes that need to be put in place. These may be related to the performance of the services, which are targeted to the users of the infrastructure, internal processes required within the PRACE infrastructure, and services required for proper integration of PRACE within the European HPC eco-system. The documents will cover the complete, but general view of the operation model, complementary to the technical specifications of the software tools and middleware functionalities, provided by the technical work packages.

¹ “Operations” is a noun whose etymology indicates: Middle English *operacioun*, from Middle French *operation*, from Latin *operation-*, *operatio*, from *operari*. This word has been traced to the 14th century

Therefore these documents should permit to give answers to the following questions:

- Which services will be provided by PRACE? (And which need to be provided outside PRACE in particular by the interconnected tier-1 and tier-2 centres)
- Which services are compulsory and provided by all PRACE installations? Which are optional and may be provided by some PRACE installations as quality or functionality enhancements?
- Which processes need to be put in place within PRACE, to deliver the services, to make the infrastructure consistent and permanent, to enable the users to use effectively the infrastructure?

The audience of these documents includes:

- End users who want to figure out what services will be provided and how to access the services.
- Stakeholders (as defined on 2.5.1 deliverable) involved in the planning and operations of tier-1 and tier-2 computing infrastructures, in order to figure out how their own services will interact with PRACE
- All technical work packages within PRACE, who need to have a reference to the definition of “what” PRACE is providing and doing in order to perform their own tasks which most often is concerned with the “how”.

We have to make the reference also to other defined operational models developed within the scope of the grid related infrastructure and initiatives. One of this is ‘*Model for Usage policy-based Resource Allocation in Grids*’ [Model for Usage Policy-based Resource Allocation in Grids, 2005, Catalin L. Dumitrescu, I Foster]. They proposed architecture and recursive model policy. They focus on brokering algorithm solving the scheduling problem of multiple jobs on multiple sites within multiple virtual organisations (VOs). They also provide some quantitative comparisons of the proposed policies. Current state of the PRACE distributed management system is that, we do not deploy multiple VOs across entire HPC ecosystem, so we think about PRACE as a single VO only.

2 Organisation of this document

The document is organised according to the following list of services:

- **Front Office** (or access to HPC) **services**, a list of services offered to the users accessing the tier-0 facilities offered by the PRACE project. These users (scientists, IT, industrial users, ...) or service providers (HPC vendors) will access the distributed PRACE infrastructures through their national tier-1 or local tier-2 centres. The services offered by the tier-0 will consist off HPC services (access on large and various kind of computational resources, access to software managed by licences), data services (through different disk spaces with different policies), pre/post processing (in order to generate their meshes or to arrange their simulations results), graphics (in order to edit their input data or to visualize their results, may be done locally or remotely).

These services are offered to grant users through different usage policies.

- **User Support and Training Services**, consist off all the services offered to PRACE tier-0 users in order to provide them support and training. These services may be considered as a part of HPC Services. A special section is dedicated in order to highlight this need and also indicate that these services may be available for non-registered users (through the PRACE Online Training Portal) or provided by third party vendors.
- **Ecosystem Integration services**, consist off all the services offered in order to provide a seamless integration of the PRACE tier-0 centres in the existing European HPC ecosystem as defined on D2.5.1. This ecosystem is formed by scientists interconnected through distributed tier-1 or tier-2 locations, by hardware and software vendors accessing for different purposes (support, pre sales, prototypes) the existing or the upcoming architectures.
- **Security services** is a transverse subset of services which ensure both data (source code, input and output files) security and integrity on PRACE HPC centres through processes of users rights isolation, data verification, SLA, facilities secure, ...
- **Back Office** (also called PRACE Office), consist off all the organisations and the internal services like system administration, operation and resources allocation policies, peer review, grant evaluation, management committee review and users committee reviews, implemented in order to offer the 4 previous services.

The sections covering these services are followed finally by more managerial aspects about management and decision support, which involve all the processes applied on order to upgrade or replace a large HPC system.

2.1 Assumptions

As the operational model definition is very dependent on the WP2 deliverables like definition of the structure of the PRACE entity, on the peer review process, ... on WP3 for the training/user support part and on the other technical WP (typically WP4, WP5, WP6, WP7 and WP8) some assumptions have to be made in the development of this deliverable.

This deliverable contains an overview gained through an extensive questionnaire of which services are offered by European HPC tier-1 generally yet involved in projects like DEISA or HPC Europa.

It will allow having a list of existing services, which will be more detailed and elaborated in D2.6.2.

The actual list of services is subject to changes during D2.6.2 depending to the discussions between partners and the results of other connected deliverables.

3 Front Office

To avoid the misunderstanding we make a note here that this section is not about HPC services but about services provided by the HPC centres on the institution-institution or human-institution level and not computing system-computing system level.

It covers human oriented interoperability in the wide sense of HPC ecosystem.

Below and for each section, we will present the survey results. For the HPC and related services every table contains the answers for the question in the first row, in the second one there are the results. The values are shown in [%] - percent of the responses for the particular question. Please note, that the values do not sum up to 100%, not all questions were answered or multiples answers were sometimes possible.

1. Is an access to HPC resources uniform per infrastructure or site?

Per infrastructure	per site	Specific
47	47	5.3

Total of 18 complete answers

2. What type of access to resources is available to users?

Command line mode	WEB mode	GUI mode	Other
100	26.3	15.3	15.3 Globus/Unicore middleware

For a total of 18 complete answers

3. Is the access restricted in some way?

through Registered subnet	through Registered machine IP	allowed from any machine
15	26	58

For a total of 18 complete answers

4. Is it possible for a user to submit, monitor and manage jobs in a uniform way common user environment, or any integrated environment? {Current state}

per infrastructure	per site	only submit per infrastructure	but not monitoring and managing
47	47	15.8	5.3

For a total of 18 complete answers

5. Should it be possible for a user to submit, monitor and manage jobs in a uniform way - common user environment, or any integrated environment? {For future Petaflop system} state}

per infrastructure	per site	Only submit per infrastructure
60	30	10

Some remarks about the wish to have a PRACE uniform submit/monitor way.

For a total of 18 complete answers

6. What kinds of jobs are available to users now?

batch mode	interactive mode	local mode	remote mode
94.4	68.4	26.2	36

For a total of 18 complete answers

7. What kind of jobs should be available to users? {in future petaflop system}

batch mode	interactive mode	local mode	remote mode
100	75	37.5	62.5

For a total of 16 complete answers and 2 skipped answers

8. Is the job checkpointing available to the users now?

system wise	User wise	no checkpointing	rely on their UPS so no checkpointing
5.6	47.4	42.1	5.8

For a total of 18 complete answers

9. Is the job checkpointing necessary in the Petaflop installation?

system wise	User wise	no checkpointing	rely on their UPS so no checkpointing
35.3	52.9	23.5	11.8

For a total of 17 complete answers and 1 skipped answer

10. Is the pending job rerouting available to the users (can users change the queue manually)?

From one queue to another on the same machine	From one system to another	No
26.3	0	73.7

For a total of 18 complete answers

11. Should the pending job rerouting be available to the users (Would users change the queue manually) in the petaflop machine?

From one queue to another on the same machine	From one system to another	No
46.7	26.7	46.7

For a total of 15 complete answers and 3 skipped answers

12. Is the resource reservation available to the users?

for the Yes answer	for the No
52.6	47.4

For a total of 18 complete answers

13. Should the resource reservation be available to the users in petaflop systems?

Yes	No
88.2	11.8

For a total of 17 complete answers and 1 skipped answer

14. Are workflow or co-allocation capabilities available to the users?

Yes	No
47.4	52.6

For a total of 18 complete answers

15. Should workflow or co-allocation capabilities be available to users in petaflop systems?

Yes	No
87.5	12.5

For a total of 16 complete answers and 2 answers skipped

16. Is a set of applications homogeneous between sites/servers?

Yes	No	Some
26.3	5.3	78.9

For a total of 18 complete answers

17. Should a set of applications be homogeneous between sites/servers in petaflop environment?

Yes	No	Some
17.6	17.6	64.7

For a total of 17 complete answers and 1 answer skipped

18. How the software licenses are managed?

Locally	Shared	Other
84.2	26.3	13.4

For a total of 18 complete answers

19. How the software licenses should be managed?

Locally	Shared	Other
71.4	50	32

For a total of 14 complete answers and 4 skipped answers

20. Do you use the federated accounting data for purposes of:

Billing	Keeping the resources thresholds	watching	No federated accounting	Other
42.1	73.7	15.8	10.5	5.5

For a total of 18 complete answers

21. Will you use the federated accounting data from petaflop systems for purposes of:

Billing	Keeping the resources thresholds	watching	No federated accounting	Other
71.4	85.7	14.3	0	28.6

For a total of 14 complete answers and 4 skipped answers

22. Are there dedicated service/services for users/groups for data upload/download and management?

Across infrastructure	Across sites	Per Server
42.1	36.8	36.8

For a total of 18 complete answers

23. Should dedicated service/services for users/groups for data upload/download and management be available on petaflop environment?

Across infrastructure	Across sites	Server
47.1	58.8	23.5

For a total of 17 complete answers and 1 skipped answers

24. Do you provide any specialized services for visualisation of the computations?

Yes	No
63.2	26.8

For a total of 18 complete answers

25. Would the provision of any specialized services for visualisation of the PETAFL0P computations be necessary?

Yes	No
75	25

For a total of 12 complete answers and 6 skipped answers

26. Are users allowed to make a backup of their data on storage systems?

Average Data size with a quite vast distribution (from 5 GB to unlimited!)	Time (from one year to infinite through HSM)	No answer
68.4	63.2	31.6

For a total of 18 complete answers

27. Should the users be allowed to make a backup of their data on storage systems in their own (in PETAFL0P environment)?

Average Data size with a quite vast distribution (from 500 GB to unlimited!)	Time (from 3 months to project life, without HSM)	No answer
64.3	64.3	28.6

For a total of 14 complete answers and 4 skipped answers

28. Is there an infrastructure dedicated for testing and development available to users? (yes; no) If yes what kind?

Yes - dedicated short queues on production systems	Yes - exclusive node(s) or partition on the production systems	Yes - other compatible machine	Yes – Interactive testing	for the No
68.4	42.1	15.8	57.9	21.1

For a total of 18 complete answers

29. Is the special infrastructure for testing and development necessary in PETAFL0P environment? (yes; no) If yes what kind?

Yes - dedicated short queues on production systems	Yes - exclusive node(s) or partition on the production systems	Yes - other compatible machine	Yes – Interactive testing	for the No
75	56.3	18.8	43.8	6.3

For a total of 16 complete answers and 2 skipped answers

30. Is there a shared workspace available for users?

Yes - it is shared between many systems in a single site	Yes - it is shared among multiple sites	for the No
83.5	5.3	10.5

For a total of 18 complete answers.

31. Should a shared workspace be available for users in the PETAFL0P environment?

Yes - it is shared between many systems in a single site	Yes - it is shared among multiple sites	for the No
75	50	0

For a total of 16 complete answers and 2 skipped answers

3.1 HPC

Many aspects of the HPC services concern all sites (computing centres), without doing the distinction on the tiers levels of ecosystem pyramid. We do not differentiate tiers (layers 0, 1, 2) in terms of the e.g. safety, data security, users' identification. The main difference between tiers is the model of sharing the resources: client-resource provider, where the resource provider is a system belonging to the tier-0 pool.

We focus on the following entities: software licenses, user accreditation, accounting and job submission.

Basing on the outcome from the D4.1.2 [11], two paragraphs present some facts regarding licensing and queues behavior.

Managing software licenses

D4.1.1 [10] states: "The schedulers need to be able to manage software licenses like any other resource, so that jobs may request access to licensed software in the job description language." Individual site software license management is dependant of the queue management availability in the batch scheduling system (D4.1.1). Since queue management is

provided by the batch scheduling system of each prototype, software licenses can be managed locally.

Question 18 outcomes shows that licenses are mainly managed locally (84%) and sometimes shared across sites (26.3%).

Software license are not planned to be managed across sites (with still 71.4% for the local management on question 19) since there are restrictions coming from licenses that do not allow cross-centre and cross-country utilisation.

By the way 50% of the respondents would like to see shared licences across sites, indicating that previous issues need to be addressed in order to increase the mutualisation.

User accreditation

Basing on peer review process users are given the rights to operate with their resources of infrastructure. The limitation of rights consists of several aspects of the utilisations:

- Evaluation of the grant proposal (peer review process)
- Computational grant requirements: number of processors, time constraints and limits, space limits
- Reporting procedure and evaluation of outcome
- Violation of rules by the users.

Current peer review practices are collected by the task 2.4 of the WP2, and analyzed in the deliverable D2.4.1 Initial Report on the Peer Review Process.

Most (but not 100% of sites) requires written agreement on the acceptance of the utilisation conditions. Of course, it does not prevent from malicious actions from the users' side. If the unaccepted action is discovered, the user loses access to their resources at once, and for the nearest future. There is no policy regarding financial responsibility for covering e.g. wasted cputime, loosing of potential cputime or sys-admin task on operating system recovery, etc. ...

For the particular answers see the questions in the Annex: "An official 'I am fully responsible for what I am doing'. Document that anyone willing to use the resources in your centre have to sign (by law)".

How the verification procedure of the new users allowed using the systems in your data centre look like?

The user presents their private ID	There is/are person that contacts all the candidates	The user presents the certified credentials by their home institution (X509-like solution)	Other with answers like
			<ul style="list-style-type: none"> • Online registration with approval devolved to project management • ssh public/private keys • Application form • Each organisation has representatives who validate the new accounts
58.8	35.3	5.9	29.4

For a total of 17 complete answers and 1 skipped

This means that every institution validates the users credential, but in different ways. In the final operational model, these procedure should be clearly stated, who and where is responsible for user credential enforcement. We have to mention, that most sites base on checking personal ID's, and 6% sites trust electronic credentials at this time. This is the

subject to be analyzed thoroughly, when Europe intends to approach towards electronic signatures and “e-country” paradigm.

Also user data operation covers many aspects, starting from the law-based features: like industrial data, medical, which are processes under external conditions formalised in the context of the particular site agreement. The regulations could impose the placement of the data – in that case, users cannot transfer files nor process jobs on another site. The final operational model will contain the guidelines on the data interoperability.

Accounting

How often do you need to report on: usage of your computing infrastructure (i.e. system utilisation) or scientific quality of computation results, other?

Monthly	Yearly	Not at all	On demand	Other with answers like:
17.6	70.6	0	41.2	20 <ul style="list-style-type: none"> • Quarterly usage statistics are provided in addition to annual report. • Every 6 months • 2x per year (I had to pick one of the answers here, to avoid error) • Each 3 months for usage of computing infrastructure and each year for scientific results

For a total of 17 complete answers and 1 skipped

Are you given the fixed report template document?

Yes	No
29.4	70.6

For a total of 17 complete answers and 1 skipped.

Accounting policy is strictly connected with the site bylaws and periodical reporting to the higher body. Every site prepares the report on the utilisation of the resources for the given timeline. Usually it is processed yearly (70%). There is no predefined template for reporting. However, some of sites prepare interim (e.g. quarterly) reports for the internal usage.

The reports are generated for the higher body, usually connected with the financing of the given site (ministry, parent institution).

There is a need to work on the common model of accounting utilisation, in particular:

- Which site is responsible for collecting the user’s activity on all granted sites?
- To whom the users should report on the used resources later on.

Question 20 of the questionnaire addressed the use of federated accounting data; in a vast majority of 73% respondent asked that this data is used for keeping the resources thresholds and for billing purposes (42%).

This second purpose is increasing to 71.4% in question 21 for the tier-0.

Quotas

There are no (space) quotas implemented at sites. In grant application form users are obliged to envisage the project space requirements, but due to the difficulty of assessment – these numbers are helpful for the planning the better space utilisation, or possible future purchase.

Job submission

In all modern batch queue systems, the user has the possibility to submit the job to a particular queue or use the workload manager to choose one automatically.

The output of question 6 shows that 94.4% of the respondents are using batch mode for submitting jobs on actual systems (and also 68.4% on interactive mode). This number grows to 100% for the output of question 7 regarding future tier-0 systems.

As shown in question 4 users are able to submit and to monitor their jobs in a uniform way common user environment.

Users for the moment are not allowed by the respondents to reroute pending jobs (only 26% on question 10 are offering this possibility) but the respondents want to see this feature expanded on future tier-0 systems with a possible rerouting on the same machine.

Also 52% of the respondents are offering resource reservation (question 12) to users, this functionality is requested by 88.2% of the respondents since tier-0 systems are considered mostly as capabilities solutions.

The batch mode approach requires knowledge about resource requirements of the job.

These can only be provided by the user, who understands the input data and can provide estimates on job parameters like total execution time, or memory and disk consumption. Modern schedulers support a “routing queue” feature, which can be used to implement single-queue submission for users.

Fulfillment of this requirement is thus more a matter of local configuration than systems software deployment.

3.2 Related services

In this subsection we will describe the services made available for users sketching out the basic functionality and the way the services may be utilized. The topics are subject of interest of WP7 as there are defined requirements for different types of the next-generation petaflop-class systems.

3.2.1 Data

The data service understood as the space made available for user's data to work on is one of the basic and mandatory services. This space however, due to technical reasons is divided into two categories – a space utilized by the running application, sometimes referenced a scratch, which is usually directly connected to the computing resource the application is running on with high performance throughput, implementing a parallel file system, with quotas enabled sometimes and not saved.

The second disk space available is used to prepare input data or analyse the results.

This second disk is usually shared among multiple machines and makes up a workspace more for the user rather than the application (used for storing personal files, sources code files or results of simulation). This space is generally managed with quotas and saved.

Because this service is one of the basic and rudimentary ones, the survey (question 30) indicated that most of the sites (close to 90%) provide this service.

Almost 83% provide shared workspace available across one site. So far only few of the sites are providing workspace shared among different sites (5%). It is unknown whether this situation is caused by technical problems or few organisations have users accessing the infrastructure from different sites. In the future however, the increase of the importance of the shared workspace is predicted. Due to the large community of the users willing to utilize the resources of petascale systems the multi-site shared workspace is a desired service for almost 50% of the respondents (question 31) and the single site workspace is needed by ~75%. There are no responses suggesting that there is no need for such service.

The PRACE deliverable D.7.5.1 [13] suggests that the size of the globally shared workspace in single tier-0 site will be 1 or more petabytes. That confirms the estimations gathered in the 2.6.1 survey.

3.2.2 Archive

An archive service could be dedicated to result files after computation or be a stand-alone service (e.g. for storing some large external data as a service) or only as an “added value” for computation. This section relates to archive service for users. The relevant questions may be: How are the time frames that the data is guaranteed to be preserved? Is there a size limit for this type of storage?

Preserving the results of the computation and the input files, often very large, on fast, shared storage is not cost effective. Therefore most of the centres are offering a service to store important data in more cost-efficient way ensuring at the same time high reliability and safety of the stored data. The answers from the survey (question 26) indicate that the archive service is sometimes confused with the backup. About 25% of the answers suggested that the backup is done either automatically by the system or by the administrators. The size of the archive storage available for the users varied from 1 to more than 500 GB (there were also cases with no limits). The retention time for the stored data differed in a great degree as well from 6 months to 10 years. At the same time the predictions regarding the future need for archive service that comes along with petaflop machines indicates that there will be great need for storage space. The predictions indicate that it is hard to foresee the required space – most of the answers (question 27) are suggesting that the available space should be customized to user's needs and the estimated need should be counted in terabytes rather than gigabytes.

3.2.3 Graphics/Visualisation

The visualisation is a topic that covers user support services and HPC services at the same time because the visualisation may require a specialized hardware. Currently about 61% offer help with visualisation to the users and (as shown on question number 24) 64% of the sites offer some visualisation services available for users in different forms (caves, curved screens etc.). At the same time in question 25 more than 75% of responses indicated that specialized visualisation services should be available for petalop systems.

Another graphics functionality is called remote graphics, it offers the possibility to perform almost all the rendering operations on the datacentre and to send to the client only compressed images in order to reduce the network bandwidth and to avoid transferring all the raw data (only pixels are transferred in order to transfer TBytes or Giga OpenGL objects is the display is just exported without any compression).

This new functionality is very interesting for large datasets (generated for example by a petaflop machine) during a study where a lot of intermediate steps for validating the data are needed.

This service has not been targeted by the survey but need to be investigated later for PRACE Front Office services.

3.2.4 Pre- and Post- processing

Some of the applications may require input in a special format, need to prepare data before computation (generation of meshes, partitioning, ...), or may produce results in format difficult to process by non-IT specialist user. The subsection is relative about all activities performed by the staff as a support in preparing the input data or converting the results to more common formats. It provides input about the support provided by the site staff for the users regarding the process of preparing the data for further computation and later data analysis, format conversion etc

Close to 68% of the respondents are yet providing support to users for preparing data and 86.7% want to see this service offered on PRACE tier-0 systems.

3.2.5 Interactive service

The process of development of new software or initial phases of computations may require access to the computing environment in an interactive way. As the contemporary computing resources are designed to be used in a batch, non-interactive way, special access scenarios or dedicated fragments of the clusters are required to provide an interactive service without disruption of the rest of the computations.

Question 28 was addressing this point by asking to the respondent about infrastructures for testing and developing purposes.

The output showed that close to 80% are offering these kinds of services to users mainly implemented with dedicated queues on production systems (68.4%), exclusive node(s) or partition on the production systems (42%) or interactive resources on frontal nodes (57.9%).

For the tier-0 systems (question 29) respondents would like to see available this service directly on production systems with dedicated queues (75%) or exclusive nodes (56.3%).

This section encompasses the procedures for providing the user with the interactive environment that is the same as the final computation environment. It may be done as a special interactive queue available for example during the day. An ssh shell interactive connection may be also considered as the interactive service. The main purpose of the service is debugging and tuning the applications for the dedicated environment. We should consider moving this section somewhere else.

3.3 Initial specifications for Front Office services

Following our discussions and the feedback obtained from the survey we are proposing this preliminary list of specifications for the Front Office services:

- Future tier-0 resources will be accessible mainly through batch submission; however interactive usage will be also possible for code development and validation or pre/post processing. This usage will be implemented with dedicated queues or exclusive resources on production systems.
- Resources reservation with workflow and co allocation capabilities will be available for users
- Because of the complexity/size of future tier-0 systems, the checkpoint / restart functionality need to be implemented by a user wise approach (more efficient but it will need information and training to users)
- Software licenses may keep managed locally
- (Remote) Visualization services need to be offered to users in order to post process results of massive simulations
- A shared workspace among multiples sites (Tier-0) is also wished in order to share personal files or source codes (\$HOME space).
- Federated accounting data will be developed in order to track billing issues or keeping the resources thresholds

4 Ecosystem integration

PRACE D2.5.1 describes the European HPC ecosystem in both academic and industrial domains.

Academic users (members of the PRACE Initiative countries or other countries accepted by the PRACE management board) once granted to an account to the futures tier-0 infrastructures would access these facilities from their tier-1 national centre or on some few cases from a tier-2 regional centre or directly from their workstation.

Industrial users (European big companies and also SMEs) would also access these facilities from their tier-1 national centre or directly from their locations.

Below we will present the results for the Ecosystem integration questions.

1. What kind of interoperation do you perform with other sites?

Common services	Common procedures	User exchange	Job Exchange	Other (in fact a no answer)
63.2	47.4	42.1	21.1	10

For a total of total of 18 complete answers

2. What kind of interoperation will you intend to perform with other sites in the petaflop environment?

Common services	Common procedures	User exchange	Job Exchange	Other (in fact a no answer)
71.4	78.6	57.1	35.7	0

For a total of total of 14 complete answers and 4 skipped answers

3. Do you interoperate with some grid infrastructure? (Yes, no)

- 71.4% for Yes with Synfiniway, EGEE framework, DEISA, globus, uncore, Chemomentum, Spanish Supercomputing Grid, and Nordic Data Grid Facility
- 78.6% for Purpose of operations with answers like access from Airbus to a BlueGene system
Offering resources to larger user communities, providing local user communities with access to other resources
Resources utilisation
Common research projects
To execute jobs on HPC machines otherwise not available to Italian HPC users
Supports complex workflow or portal interface to standard packages
Distributed Nordic tier-1
Resource exchange, benchmarking; LHC research
For experimental purpose
- 21.1% for the No answer

For a total of 18 complete answers

4. Can you see the advantage of interoperation petaflop systems with some grid infrastructure?

- 66.7% for Yes with answers like Deisa, globus, Chemomentum, D-Grid and EGI
- 60% for the Purpose of Operations with answers like
Workflows (locally pre-process, tier-0 compute, locally or tier-0 post-process)
Remote workflow, access to chemical (toxicity) databases
To implement the PRACE vision (pyramid of tier-0 and tier-1)
Requirement of some (but not all) user groups
Easy migration of jobs and data between tier-0 and tier-1 systems
Common user environment
Submitting jobs for farming applications
AAA, data movers, current DEISA users are the most potential petaflop users

For a total of 15 complete answers and 3 skipped

Another potential community, which need to be considered, consist off third party actors like hardware and software vendors or HPC service providers. In this case access to future PRACE tier0 infrastructures will be relevant for

- R&D activities in collaboration with tier0 centres or PRACE connected groups like STRATOS
- Prototypes
- Installation phases during procurements
- Potential user support or tier-0 support after procurements

D4.2.1 describes existing solutions for ecosystem integration, which are used in several distributed infrastructures. The deliverable describes the requirements for the integration of a tier-0 infrastructure in an ecosystem that includes other tier-0s, tier-1s and other research infrastructures or potential industrial users. The deliverable builds on deliverable D4.1.1 that includes a requirements analysis for tier-0 system management and D4.1.2 that reports on existing solutions for tier-0 system management.

In this section, we summarize some of the responses from the questionnaire regarding ecosystem integration as well as some of the main findings of deliverables of D4.1.1, D4.1.2 and D4.2.1 [12] that are relevant when considering operational requirements and services for a tier-0 infrastructure in an ecosystem that includes multiple tier-0, tier-1, and other research infrastructures or potential industrial users.

In question 1 we saw that for all the respondents which yet are interconnected with other sites the main reasons are to implement common services and procedure or to exchange job or users.

For the future tier-0 systems the respondents asked for more common procedures between interconnected sites.

As shown on question 3 a total of 71% of the respondents of the questionnaire answered that their tier-1 national centre interoperates yet in some way with an existing distributed/grid infrastructure.

These infrastructures include e.g., EGEE, DEISA (close to the half of the respondents are DEISA members), the Spanish Supercomputing Grid, Synfiniway, Chemomentum, and the Nordic Data Grid Facility (NDGF). Technologies used for the current interoperation include UNICORE, Globus, gLite and ARC. The interoperation has been established for various reasons, e.g., to provide access to other resources (e.g., specialized hardware), to support complex workflow or portal interfaces to standard packages, and for benchmarking and experimental purposes.

Many (67%) of the respondents in question 4 see the advantage of interoperation between the future tier-0 systems with existing grid infrastructures. The purpose for this interoperation is manifold: to implement the European HPC performance pyramid model, to enable distributed workflows (e.g., locally pre- or post-process, access to remote databases), migrate jobs and data between tier-0 and tier-1 systems and other infrastructures, and establish common user environments across the entire ecosystem.

The following areas are addressed in D4.1.1, D4.1.2 and D4.2.1:

1. Network connectivity ensures that the infrastructure of geographically distributed sites is interconnected. The sheer size of the input and output data sets that can be expected for many applications on a tier-0 system, requires high bandwidth data transfer to/from the tier-0

infrastructure. Some of the transfers may also be time-critical and should not be influenced by other traffic that may exist on a shared network connection. This necessitates the existence of dedicated network links between the resources (incl. other tier-0s) that are included in the distributed infrastructure that constitutes a scientific application. To enable high-speed data transfer, the end-to-end network performance needs to be monitored and optimized, network topology needs to be optimized, and (compatible) network addressing ensured.

To integrate the PRACE infrastructure with other parts of the ecosystem, the GÉANT network will be used. The PRACE prototype systems will be connected to this academy network provider, which will ensure their connection to the ecosystem.

In order to address potential industrial needs, discussions with GÉANT (at the European level) and also with national network providers need to be opened to allow industrial usage of academics networks.

2. Authentication, Authorisation and Accounting (AAA) define in short how users are identified, what resources users are allowed to access (and use), and how users are charged for resource consumption? This requires the establishment of uniform or interoperable mechanisms for user identification, mechanisms for granting access to resources and tools for logging resource usage. Also the homogeneity of the user (and project) name space across the infrastructure (or even total ecosystem) must be ensured.

In deliverable D4.1.1 it was stated that a mandatory requirement is to provide a uniform way to access all the tier-0 systems. A single-sign-on mechanism is needed to allow users to be authenticated and identified across all sites. Deliverable D4.2.1 describes that X.509 certificates, already provided to the users, would be accepted to authenticate on tier-0 sites.

Authorisation and accounting are not mandatory topics for integration of tier-0 systems as these concern information and data that are specific to a given infrastructure. There is not necessarily an integration gain by exchanging them.

3. Data management concerns the handling of data required or produced by computations and enabling users to move data across the infrastructure (e.g., between tier-0 and tier-1 sites). This requires the definition of data transfer protocols, wide-area access protocols, and metadata management.

Are there dedicated service/services for users/groups for data upload/download and management?

Across infrastructure	Across sites	Per Server
42.1	36.8	36.8

For a total of 18 complete answers

Should dedicated service/services for users/groups for data upload/download and management be available on petaflop environment?

Across infrastructure	Across sites	Server
47.1	58.8	23.5

For a total of 17 complete answers and 1 skipped answers

The questionnaire showed that a big majority of respondent are offering services for data upload/download and they want to see this services offered (and expended across sites) for future tier-0 systems.

A common interface and protocol is needed in order to be able to move data between tier-0 and tier-1 systems and other infrastructures. At present, GridFTP provides a de facto lingua franca understood by a large range of storage systems. It is supported on many infrastructures.

D4.1.2 states that GridFTP will be evaluated as the basis for transferring files between tier-0 sites; it is advisable that GridFTP interfaces are exposed by tier-0s towards the rest of the ecosystem in order to ensure maximum compatibility. It is up to the individual sites to decide which GridFTP server solution best suits the particular storage systems that are in use.

For proper integration into the ecosystem, a file transfer service that is capable of unattended managing large file transfers between two GridFTP endpoints is desirable. Prime candidates for evaluation are the EGEE File Transfer Service (FTS), the Globus Reliable File Transfer (RFT), SDSC's Storage Resource Broker (SRB) and its successor iRODS, SIMDAT, and dCache.

Solutions that will emerge from the scouting and development done in DEISA2 WP4 task "Data-related technologies" will of course be considered as well.

According to D4.1.1, "real-time sharing of data between tier-0 sites is not a mandatory requirement, i.e., it is not mandatory to provide a distributed shared file system".

Stage-in/out features are essential for data-processing jobs; support for staging files from tier-0 storage into a job working disk space is definitely a required service. It is also highly desirable that the stage-in/out tasks can be integrated into the batch scheduler's job execution workflow, so that computing time is not wasted when waiting for data to be transferred. DEISA already provides software support for these features, integrated with major batch scheduling systems. The DEISA solution is a prime candidate for installation at PRACE tier-0 sites; only actual usage will tell if there is any deficiency or gap to be filled.

4. Monitoring of resources. A set of monitoring tools is needed to check system health and notify users about the status and possible malfunctioning in the infrastructure. All layers in the infrastructure need to be monitored: hardware, operating system, network, middleware and application.

The exchange of monitoring information between tier-0 centres and the ecosystem is not a mandatory requirement. However, it would be interesting for tier-0 sites to be able to publish monitoring information to the relevant actors in the HPC ecosystem. D.4.1.2 highlights that an important requirement is the ability for tier-0 sites to be able to expose monitoring information with a fine-grained access web service interface.

5. Resource management and allocation. The available computational power of a tier-0 system must be shared in a fair (or predefined) manner between users. This includes remote resource allocation (or advance reservation), workflow capabilities, and possibly license management.

In most of its aspects, resource management and allocation is however not a relevant topic for ecosystem integration. On the other hand, according to D4.1.1, mapping to a standardized job description language is mandatory in order to fulfil the vision of users being able to submit jobs to the whole tier-0 infrastructure using a uniform interface. To accomplish this, each tier-0 local scheduler has to be able to either directly use a standard job description language or translate this to a locally used specific job submission language.

Features for job management systems like job checkpoint and restart, job rerouting and co-allocation, advance reservation, backfill functionalities and programmable API allowing to inquire about batch system parameters often only make sense for tier-0/tier-1 integration with identical platforms and local resources manager. PRACE WP4 will rely on DEISA to provide

these capabilities were supported by the local site, and only in view of a tier-0/tier-1 integration between similar platforms.

6. Access to the resources to enable users to perform their computations. This can be through grid middleware, command-line access, web portals and GUIs.

What type of access to resources is available to users?

Command line mode	WEB mode	GUI mode	Other
100	26.3	15.3	15.3 Globus/Unicore middleware

For a total of 18 complete answers

The integration of the resource access mechanisms of PRACE in the ecosystem is done by providing the same access tools as in existing infrastructures.

Interactive command-line based access to machines may be provided through a secure ssh service compatible with GSI (Grid Security Infrastructure) like GSI-SSH.

Users can install clients on their workstations and access to sites which allow a GSI-SSH kind of access to users. tier-0 sites will have UNICORE installed, most likely version 6. UNICORE provides a graphical user interface that can be installed on the end user's workstation, as well as a command-line based client (DESHL, a standard based Access to an Heterogeneous European Supercomputing Infrastructure). To interface with other infrastructures, Globus WS-GRAM (Globus Resource Allocation and Management) should be provided on the tier-0 systems.

Web portal interfaces and its application to tier-0 systems needs to be discussed in PRACE. One technical possibility is to use portal implementation that is being developed in DEISA. Portal jobs are handled via the UNICORE infrastructure.

7. Application and user support: providing tools that are required for the computations and providing help to the users. For a distributed infrastructure, this includes a common production environment to enable data, users and applications to move easily between systems in the infrastructure, and a help desk functionality. This is addressed in more detail in Sections 3 and 5.

It is not relevant for PRACE to integrate the user support with other infrastructures such as EGEE/EGI, since the technical teams of PRACE and EGEE/EGI will not work in close cooperation with each other. The situation is different for DEISA, whose majority of partners is also PRACE partner. For PRACE and DEISA, it is relevant to have common user support functionalities and share a common software environment, or a subset of it. It is natural that PRACE adopts the technologies and tools used in DEISA.

8. Initial specifications for the Ecosystem Insertion services: Following our discussions and the feedback obtained from the survey we are proposing this preliminary list of specifications for the Ecosystem insertion services:

- A lot of PRACE's partners tier-1 systems are interconnected with existing grid infrastructures (like DEISA, EGEE, NDGF, ...). Future tier-0 systems will need to facilitate interoperation with these structures in order to implement common procedures and services.
- An uniform way to access all the tier-0 systems through single sign-on

- There is a need to implement a common interface for data movement between tier-0 and tier-1 systems and other infrastructures
- Future tier-0 systems will be accessible through Command Line interface, web access, GUI or existing middleware like Globus/Unicore

5 Training, support and dissemination

5.1 User support

User support services in existing HPC centres are organized in several different ways, but the most common is the central point of requests with internal workflow of responsibilities (~78% of responses). However, in some centres the user requests may (or need to) be sent directly to administrators of specific resources. In some cases, if there are both ways available, the user can take his preferred option.

Below we will present the results for the questionnaire's questions relatives to User support.

1. What kinds of interfaces are available to users to request for support?

Web	Email	Phone	Service Desk
61.1	100	83.3	44.4

For a total of 18 complete answers

2. What is the workflow and responsibilities to solve the users' problems? (i.e. requests are forwarded from the support team for the infrastructure to the site administrators, then to a system or application administrator; to other users; other)

Central point of requests	Direct requests to the administrator of (system, application, network, security,...)
77.8	38.9

For a total of 18 complete answers

3. Are the user support processes certified or implemented in accordance with standards? (i.e. ISO 20000, or other) If yes, which one?

Yes answers with ISO9001 or ISO27000 or ITIL	No
27.8	72.2

For a total of 18 complete answers

4. Do you think, that the user support should be implemented in accordance with standards in petaclass environment? (i.e. ISO 20000, or other) If yes, which one?

Yes answers with ISO9001 or ISO27000 or ITIL	No
33.3	66.7

For a total of 12 complete answers and 6 skipped

5. What user support services are available now?

Help with preparation of data for computation	Help with post-processing, i.e. visualisation	Parallelisation of application code (from sequential handed)	Help with tuning of parallelized application	Help with tuning of any application code	Help with algorithms	Help with numerics	Help with porting applications	Other
66.7	61.1	61.1	94.4	77.8	33.3	33.3	94.4	6

For a total of 18 complete answers

6. What user support services should be available in the future petaclass environment?

Help with preparation of data for computation	Help with post-processing, i.e. visualisation	Parallelisation of application code (from sequential handed)	Help with tuning of parallelized application	Help with tuning of any application code	Help with algorithms	Help with numerics	Help with porting applications
86.7	80	80	100	66.7	80	73.3	93.3

For a total of 15 complete answers and 3 skipped

Question 1 and 2 answers shows that the user support service is available through several centralized channels (starting from the most common): email (100%), phone call (~83%), Webpage (~62%), or at Service Desk (~44%). Usually the subset of these communication channels is used.

As answered in question 3 the processes of user support in several centres are certified or implemented in accordance with standards (~28%), but most of the centres don't certify the processes (~70%). In some cases the ISO 9001, ISO 27000, ISO 20000 and ITIL (Information Technology Infrastructure Library) are implemented.

ISO 9001 is one of the ISO 9000 family of standards maintained by ISO, the International Organisation for Standardisation, and is administered by accreditation and certification bodies. The standard is intended for use in any organisation which provides any form of service. It defines a number of requirements which an organisation needs to fulfill if it is to achieve customer satisfaction through consistent products and services which meet customer expectations.

The ISO 27000 series provides best practice recommendations on information security management, risks and controls. It is similar in design to management systems for quality assurance (the ISO 9000 series). The ISO 20000 standard comprises two parts: a specification for IT Service Management and a code of practice for service management.

The ITIL is a set of concepts and policies for managing information technology (IT) infrastructure, development and operations.

Applying one of the above standards (or any other appropriate one) to the processes of user support in future PRACE RI was supported by more than 33% of responders of question 4. In that case the ISO 9001 or ISO 20000 standards were recommended.

Currently available support services directed to HPC users include:

- Help with tuning of applications (parallelized and not) – most of the Centres.
- Help with porting applications – most of the Centres.
- Help with preparation of data for computation.
- Parallelisation of the application code.
- Help with post-processing (i.e. visualisation).

Some of the centres help users also with algorithms and numeric libraries, or offer additional science support regarding e.g. recommendations of software packages appropriate for the a research problem.

As shown in question 6 the responders recommended all above support areas to be included into the future PRACE RI support service, but three of them: “Help with preparation of data for computation”, “Help with tuning of parallel applications”, and “Help with porting applications” are considered as most important (100%). The second choice is “Parallelisation of the application code”, “Help with tuning of parallel applications”, and “Help with porting applications” (~92%). This requires an agreement between tier-0 sites on dedicating personnel to these services and distribution of the responsibility.

Possible models for user support service operation regarding troubleshooting include:

- A global ticketing system that links together regional ticketing systems.
- Separate ticketing system at every tier-0 site.

5.2 User training

The Work Package 3 is responsible for developing and implementation of a European HPC education and training program for scalable scientific computing. As stated out in deliverable D3.1.1 “Final plan for the use and dissemination of foreground”, this training program will be continued within the permanent PRACE Research Infrastructure. Thus, there should be prepared a plan for user training, based on responses on a survey of training needs and requirements (D3.3.1) and conclusions from the Summer School (D3.3.2) and upcoming Winter School (D3.3.3). Answers on some additional questions collected during the survey of 2.6.1 task give a general view of current organisation processes and concepts.

Below we will present the results for the questionnaire’s questions relatives to User training.

1. What kind of training is available to users?

Group	Personal	Periodical	On Demand/request	Web training material
76.5	23.5	41.2	52.9	6

For a total of 17 complete answers and 1 skipped

2. What kind of training should be available to users in the petaclass environment?

Group	Personal	Periodical	On Demand/request	Web training material
92.9	14.3	64.3	64.3	6

For a total of 14 complete answers and 4 skipped

3. Are the courses available only to the site’s users or are open for others?

Open-Free answer	Open-Higher fee answer	Only site registered people
92.9	14.3	11.8

For a total of 17 complete answers and 1 skipped

4. Should the courses be available only to the site’s users or left open for others in the future?

Open-Free answer	Open-Higher fee answer	Only site registered people
86.7	20	6.7

For a total of 15 complete answers and 3 skipped

5. Are your hardware/system/application vendors performing training for the users?

Yes	No	Other
82.4	23.5	23.5 : <ul style="list-style-type: none"> ○ One or two initial courses are provided in the first year of operation ○ Yes, but the effort spent on this is limited, ○ Depends on needs! ○ Additional training by vendors

For a total of 17 complete answers and 1 skipped

6. Should your hardware/system/application vendors perform training for the users in the future?

Yes	No
85.7	14.3

For a total of 14 complete answers and 4 skipped

7. Do you have a training program for your personnel?

Yes	No
47.1	52.9

For a total of 17 complete answers and 1 skipped

8. Are you willing to have a training program for your personnel in the future petaflop environment?

Yes	No
93.3	6.7

For a total of 15 complete answers and 3 skipped

9. Provide a list of courses that are available at your site (or URL to the list available on-line) Titles or URL

12 complete answers with:

- We don't have any list on our site
- We have an annual MPI/OpenMP course given by Ralf Rabenseifner (HLRS), and we are currently preparing our full course program of other material.
- <http://www.pdc.kth.se/events/events2008/view>
- Courses are prepared on demand
- <http://www.icm.edu.pl/kdm/Szkolenia>
- <http://www.cineca.it/scuoleestive/>
- <http://www.hpcx.ac.uk/support/training/index.html>
- <http://www.hector.ac.uk/support/cse/schedule/>
- <http://www.lrz-muenchen.de/services/compute/courses/>
- <http://szkolenia.man.poznan.pl>
- The list is available on the CCRT web site

Question1 outcome shows that the current training services focus on courses for groups of users (in a few cases personal training is performed), which are organized mostly on demand, but some of them are periodical. In one case the on-line training material is used. This is also the recommendation for the future PRACE training service: courses for groups (93%), personal (~14%), organized on demand (~64%), with some of them (e.g. basics, most wanted) periodical (~64%), and some performed on-line.

The courses are mostly available to open audience and free of charge (~93% in question 3). In some cases the course is free for site registered users or academic community, with an option for others to participate (with some fee, sometimes higher for industrial organisations). For the future service open and free of charge courses are mostly recommended (~87% in question 4). Some responders (~7%) suggested that courses dedicated for registered users or academic community should also be available, with small fees for others. This could be reasonable in the case of courses that are very specific, expensive or restricted in some way due to license issues.

The PRACE training program for potential and existing users should contain courses regarding High-End Computing. It could consist of several courses organized by PRACE RI and dedicated to PRACE tier-0 sites' usage, but some general courses on parallel programming methods organized by other sites, could also be included. The training program should be complementary with the support services, as it is the way to educate and enhance the competences of the users and their awareness of new HPC methods and architectures.

Possible areas of the training include:

- The usage of petascale systems available within PRACE tier-0 sites (e.g. job submission, security issues).
- Porting of users' applications' to new deployed and planned architectures.
- Basics of Multi-Core Programming, Parallel Programming Models.
- Basics of Programming Accelerators/GPUs, PGAS Languages.
- Application specific courses (e.g. functionalities, possible appliance to the research area, usage strategies).

Some of the courses, especially regarding concrete applications, is currently prepared and performed in cooperation with vendors (~83% as reflected on question 5). The support for this kind of training is about 83%. This is the additional way to transfer the knowledge from vendors to users, if the centre (or future PRACE RI) doesn't have a specialist in the subject among its personnel. Although the application training could be performed at any tier-0 site (e.g. the application is only available on that site), it seems to be reasonable to coordinate the whole process at the PRACE RI level.

The PRACE RI should maintain a single point of information about courses available for PRACE users. This could be done by a section of the PRACE RI Website or a specific training portal. The training materials and exercises could be also available on-line for attendees of courses. The on-line training techniques should also be explored, to give the PRACE users possibility of getting a basic knowledge or further study on a specific subject.

5.3 Personnel training

Question 7 outcome shows that about 47% of the HPC centres have a training program for its personnel implemented. This includes courses about: the deployed hardware and software (operation systems, middleware, libraries, and applications), application tuning and development, application porting to new architectures. Some of the support services require advanced knowledge in programming methods and languages, system and middleware configuration and adoption, and others.

In question 8 close to 93% of the responders support the idea of a personnel-training program for the PRACE RI. Thus the future PRACE RI should take into account the constant need for improving qualifications of the personnel, especially in the case of deploying new architectures. Some of the training could be held in cooperation with vendors as a part of the procurement contracts.

5.4 Dissemination and outreach

The analysis of the survey and WP3 deliverables shows that there is a need to define the dissemination and outreach process and activities to be performed by PRACE RI after the end of the PRACE project. The dissemination actions are widely performed by the HPC centres. Thus, one of the tasks during operation of the PRACE RI should be the dissemination about the organisation itself, its services and usage.

Below we will present the results for the questionnaire's questions relatives to Dissemination and Outreach.

1. What kind of dissemination actions do you perform regarding your infrastructure or site and its usage? Provide a list of activities.

Conference papers	Stands/Exhibitions	International/local events about computing	User forums	Journals	Printed Newsletters	Advertising	Other
52.9	52.9	76.5	58.8	29.4	35.3	5.9	12

Other: no dissemination actions and Electronic Newsletters.

For a total of 17 complete answers and 1 skipped

2. Do you have dedicated personnel to perform dissemination and outreach?

Yes	No
58.8	41.2

For a total of 17 complete answers and 1 skipped

The role of the PRACE RI dissemination activities would be to keep all HPC partners within the member countries, international bodies, industry and the public at large aware of PRACE RI, its services and usage in research, and the strategic importance of High Performance Computing for Europe. This requires constant communication and interaction with industry partners, decision makers and funding bodies.

The target groups and dissemination channels for the PRACE project are identified in the Deliverable 3.1.1 “Final plan for the use and dissemination of foreground”. The permanent PRACE RI should continue the activities for the identified groups:

- Scientific communities as potential users.
- Industry as potential users, collaboration partners and vendors.
- Policy makers and funding agencies.
- Research infrastructures, organisations and the general public.

In question 1 the main identified areas of the dissemination activities include:

- Website – as a continuation of the PRACE project website (<http://www.prace-project.eu>). The editorial team should be established in order to maintain and update the site with information about the mission, resources, activities and achievements of PRACE RI within the European HPC Ecosystem. The additional features include: an RSS feed system for readers, a visitor statistics monitoring system.
- Newsletter – as a continuation of the PRACE project newsletter. It should be available also from the website, with possibility to subscribe/unsubscribe. The newsletter could be targeted to HPC stakeholders, media, industry, research and scientific communities and political decision makers. It should be released periodically (e.g. quarterly). Some of the centres publish also printed versions of the newsletters.
- Dissemination media pack – this includes: general brochure, flyer, poster, press releases, and give-aways.
- Conference papers
- Attendance at international or local events about computing, including user forums, and organizing booths/exhibitions.
- Articles in journals (e.g. about the usage and most important results of computation, success stories).

None of the HPC centres advertise itself in the manner of commercial advertisement.

Most of the centres have dedicated personnel to perform dissemination and outreach (~59%), and this is an early recommendation for the future PRACE RI. The activities should be performed on both the global (PRACE RI) and the local (HPC centres) level, depending on the final distribution of the responsibility between the PRACE Partners.

5.5 Vendor Documentation & training with regard to vendor tools

The transfer of the knowledge from the vendors to PRACE RI personnel or users is usually done on two levels:

- Passive: through documentation delivered with solutions.
- Active: through a training program regarding deployed solutions.

Some of the centres declare that vendors are performing training for the users. That kind of training seems to be important with regard to specific vendor tools deployed on the systems. Thus, if such tools will be deployed on PRACE environment, they could be included into a training program for the PRACE users. The PRACE personnel could also be trained, to improve the quality of support services.

5.6 Initial specifications for the Training, Support and Dissemination services

Following our discussions and the feedback obtained from the survey we are proposing this preliminary list of specifications for the Training, Support and Dissemination services:

- User support may continue to be addressed through central points of request with different channels like email, phone, web or service desk since email and phone may be used preferably
- There is no immediate need to accordance to standards like ISO or ITIL but this kind of certifications may be asked by some potential communities like industry
- Due to the kind/complexity/size of future tier-0 systems user support services will need to focus on
 - Parallelization and tuning of parallel applications
 - Pre and post processing (visualization)
 - Help on Algorithmic and numerics
- Courses will be preferably organized periodically for user groups and on-line training will be available, courses will be mostly available to open audience and free of charge
- The training program should be complementary with the support services
- A permanent training program for tier-0 personnel should be available especially in the case of deploying new architectures. Some of the training could be held in cooperation with vendors as a part of the procurement contracts.
- In general more involvements of HPC vendors is asked for training issues

6 Security

The scope of the security subject is very broad as this is one of the key factors in the IT world. In the context of the technical operations of the PRACE project the scope narrows to the following issues:

- Legal requirements and responsibilities
- Other constraints (ownership and confidentiality of data, source code, ...)
- Definition of a common security level across the distributed infrastructure in order to offer a minimum and common level of security services
- Concepts and organisation² for
 - Site security
 - Data security

The legal requirements regarding certain aspects of the operation model are regulated by the law of the country where given institution resides. In this chapter however we will focus on the mutual obligations between the HPC service provider and the service consumer as this is regulated by the agreements.

Below we will present the results for the questionnaire's questions relatives to Security services.

1. Is there a uniform user authentication and authorisation service available?

Across the Infrastructure	Across the site	No	Other
35.3	58.8	11.8	5.9

For a total of 17 complete answers and 1 skipped

2. Do you allow your users to submit/run their own programs?

If yes, are the applications verified on the level of code, scanned for malicious content etc?

Yes	If yes, are the applications verified on the level of code, scanned for malicious content, etc?	No
94.1	0	23.5

For a total of 17 complete answers and 1 skipped

3. Do you have user activity-tracking software on your systems?

Yes, always on	Yes, on demand is on	No
29.4	35.3	35.3

For a total of 17 complete answers and 1 skipped

4. Do you provide your users with an SLA regarding the uptime, confidentiality, data consistency etc?

Uptime	Confidentiality	Data consistency	Other
35.3	29.4	29.4	52.9 <ul style="list-style-type: none"> • No SLA • Only requested from commercial users • Best effort • Agreement for service downtime twice a month: 1.&3.Thursday. • Don't know of any

For a total of 17 complete answers and 1 skipped

² The technical aspects do not belong to this document, if needed; we should have a corresponding technical document.

5. Are your users responsible for any possible damage caused by their behavior, i.e.: security breach because of trivial password or running malicious software, or giving the credentials to someone else? (yes; no) If yes, what are the possible consequences?

Yes with answers like:	No
82.4 <ul style="list-style-type: none"> • Giving credentials to someone else results in a life ban from using the Centre's facilities (reminding people of this fact acts as a good deterrent). • Block users; there might of course also be legal consequences. • Closing an account • Consequences are those listed by the Italian laws. From administrative computer crime to criminal computer crime. • Loss of access • Closing of account, termination of access • No further access to the system • Revocation of access 	5.9

For a total of 17 complete answers and 1 skipped

6. What procedures do you implement to prevent and limit undesired user behavior?

On-line training for users about basic security issues	an official “I am fully responsible for what I am doing”. Document that anyone willing to use the resources in your centre have to sign (bylaw).	Other with answers like: <ul style="list-style-type: none"> ○ Activity scans ○ None really ○ Acceptance of the Rules and Regulation with application ○ Guide lines on www and application form
17.6	82.4	29.4

For a total of 17 complete answers and 1 skipped

Apart from the responsibility issues there may be also other subjects tightly coupled with the term security. Beginning with the very basic problems e.g. how is the verification of new users performed, are the applications submitted by users checked in any way etc.

The view how the security as a level of service for the users is perceived is quite coherent in all sites. In most cases the institution has its own server room either in shared-purpose building or dedicated building. All the computer rooms are protected by different arts of identification mechanisms so the physical security of the systems matters. Additionally in some cases there is a staff dedicated for this purpose.

The verification procedure of the new users that should be given access to differs significantly between the sites. In general users are required to register somehow by filling a hard copy or on-line form. Just after in question 11 the outputs show that around 30% of the sites have people that personally contact the users. In about 59% cases users are required to present personal ID numbers as shown in question 11 outputs.

7. Do you have your own data centre room?

The answer was 100% for the Yes answer, for a total of 17 complete answers and 1 skipped

8. Is your data centre located in a dedicated building or is it a shared-purpose building (offices, other companies etc)

Dedicated Building	Shared purpose building
35.3	64.7

For a total of 17 complete answers and 1 skipped.

9. What are the physical access restrictions, physical means for preventing unlawful access?

The 17 complete answers (1 skipped) were the following:

- Badge system
- Standard office security - locked doors, electronic key access. The machine room is protected by strong security doors.
- Door access control, movement sensors
- Token authorisation system
- Electronic card + biometrics
- Guards, special locks, CCTV, alarms,
- Electronic card controlled access to several building rooms
- Details restricted
- Key card / pin code access
- Access keys, doors, video, personnel
- This is a protected site with physical access restrictions
- Security guards that control the access to the building, plus TV monitoring, alarms, etc
- Only authorized personnel, proximity cards, biometric access system

10. What is your policy of granting the access to the data centre room to non-staff persons (e.g. visitors, vendors, facility workers, cleaning personnel)?

The 17 complete answers (1 skipped) were the following:

- Badge access restrictions
- Access to the machine room for non-authorized personnel (including staff without machine-room access) requires the machine room doors to be opened by the facility management personnel.
- Decided on demand
- Access only with the personnel
- Passport ID is required
- On request policy, implemented by means of electronic card controlled access
- Case by case. Visitors and one-time workers are always accompanied by staff
- Visitors, vendors, facility workers only accompanied by staff, cleaning personnel has their own set of keys
- We have limited access to the data centre room through badge's reader

11. How the verification procedure of the new users allowed using the systems in your data centre look like?

The user presents their private ID	There is/are person that contacts all the candidates	The user presents the certified credentials by their home institution (X509-like solution)	Other with answers like
58.8	35.3	5.9	<ul style="list-style-type: none"> • Online registration with approval devolved to project management • ssh public/private keys • Application form • Each organisation has representatives who validate the new accounts
			29.4

For a total of 17 complete answers and 1 skipped

As show in question 1, most of the sites (~88%) have unified mechanisms either covering entire site (58%) or infrastructure (~35%). In some cases there are other models of authorisation such as identification per user class or type of project (e.g. astrophysical, chemistry, biology,).

Generally, once authorized, the users are trusted when it goes to what they are doing on the machines. As exposed by question 3 only ~ 35% sites is monitoring user activity on the servers. In all cases users are allowed to run customized or self-written code without any form of code verification or binary checking. In most cases the users are responsible for any damages their software may cause (more than 50% of relevant answers in question 4). The consequences for the users that will not obey the terms of use and cause a security breach or damage is not severe; the only punishment is revocation of access to the computing infrastructure.

In question 4 most of the centres give some sort of Service Level Agreement for users willing to use its resources, nevertheless there are cases (less than 20%) where no guaranties are agreed. The scope of the guarantied services differs significantly – some the centres (more than 35%) promise to keep the uptime of the servers on given level, some guarantee data consistency (~ 30%). What is interesting – only about 30% promises to keep confidentiality of the information. Does it mean that this is irrelevant for the users or perhaps the providers are not willing to give such guaranties? On the other hand the current survey gave no hints regarding the consequences, legal, financial or administrative, for the supercomputing centres that are not fulfilling the terms included in the SLA.

Is the job checkpointing available to the users now?

system wise	user wise	no checkpointing	rely on their UPS so no checkpointing
5.6	47.4	42.1	5.8

For a total of 18 complete answers

Is the job checkpointing necessary in the Petaflop installation?

system wise	user wise	no checkpointing	rely on their UPS so no checkpointing
35.3	52.9	23.5	11.8

For a total of 17 complete answers and 1 skipped answers

It looks that fault-tolerance is perceived as a non-existent or not important problem of the users because on one hand the SLA not always covers the uptimes and on the other hand other fault-tolerance mechanisms such as checkpointing are considered either as unnecessary (more than 42%) or are treated as a problem the user has to deal with (47%). For the moment only 5.6% relies on system wise checkpoint restart facilities mainly due to a lack of solutions provided by vendors.

What is interesting the checkpointing, perhaps due to the sheer size of the future systems, is perceived as a feature that should be implemented either as a part of the system (35% which represents a challenge for the vendors or the HPC community) or as part of user application (52%). Only less than 12% answers indicated that such feature is not required because of e.g. redundant power supplies.

Taking this into account the fact that the providers of the computing power are very reluctant to promise anything to users on one hand, and lack of real consequences for any agreement breaches done by the users on the other, it is hard to speak of any general legal requirements and responsibilities. Probably in case of real losses it is up to any individual organisation or person to perform any legal actions to get the compensations if a SLA has been firmed between the provider of the service and the user.

The data of the users is protected either by automatic backups or the users are free to do backups on service made exposed by the site administrators. This is interesting in scope of SLA for the users where only 20% of the sites guarantied the data consistency, and at the same time backs up the data in either automatic or manual way.

6.1 Initial specifications for the Security services

Following our discussions and the feedback obtained from the survey we are proposing this preliminary list of specifications for the Security services:

- All the institutions have their own server room either in shared-purpose building or dedicated building. All the computer rooms are protected by different arts of identification mechanisms so the physical security of the systems matters.
- Few SLA features are offered for the moment by the centres, this point deserves to be discussed in deep in 2.6.2

7 Back Office

The PRACE Service is similar to a Back Office and describes all the organisation's parts and the internal services like system administration, operation and resources allocation policies, peer review, grant evaluation, management committee review and users' committee reviews. As the PRACE RI is considered to be a single entity, it remains reasonable to prepare a set of common processes and procedures to implement common internal services. At the same time the implementation of processes may differ in details between partners, as it may relay on the existing processes of the site.

Resource usage policy issues can arise within infrastructures that integrate participants and resources spanning multiple distinct physical sites. Participants may wish to delegate to PRACE RI the right to use certain resources subject to local policy and service level agreements.

We propose policy model, and define roles and functions, for scheduling resources in PRACE RI while satisfying local resource owners and corresponding policies tackled in Memorandum of Understanding, PRACE related agreements.

7.1 Handling of exceptional resource requests

How are such local and PRACE RI policies to be expressed, discovered, interpreted, and enforced? As a first step to addressing these questions, we gathered basic data regarding job related paradigms via designed survey (see the Appendix). Remaining knowledge has come from the particular site representatives. We make the assessment of the local resource allocation policies, if any, without any specialized context. We also provide the initial specification of the policy to be considered in PRACE RI but without going into details, which is the starting point for the next deliverable ready for M24.

In many organisational units ones encountered problems in case of processing occasional requests. In this section we would like to summarise most of them, presenting the current state of the servicing of toll-orders.

All computing centres have the regulations regarding the resource utilisation by the user that should be signed by them. However, they are very common, and not tackle the rare situations.

1. How do you handle very LARGE jobs?

Do the users need to consult the site's Council of Users (higher body) every time before running the job?	Do the users need to consult the site's administrators every time before running the job?	Are there grants with special privileges i.e. additional disk space, or better Priority?	Dedicated queues	Very large jobs are not supported here.	Other
5.9	17.6	35.3	76.5	11.8	29.4 : <ul style="list-style-type: none"> • Anyone can run a large job without special permission. Part of the scheduling policy is to weight jobs according to job size in addition to other priority factors. • Not clear what very large is - our users can use up to the full capacity of a certain machine • Queue structures usually give priority to jobs (e.g. large jobs) that cannot be run at smaller facilities. • Prioritisation of large jobs • If the users need more than half of the computing power, they have to ask the agreement of the management committee

For a total of 17 complete answers and 1 skipped

2. How do you handle very LONG jobs?

Do the users need to consult the site's Council of Users (higher body) every time before running the job?	Do the users need to consult the site's administrators every time before running the job?	Are there grants with special privileges i.e. additional disk space, or better Priority? Dedicated queues.	Dedicated queues	Very long jobs are not supported here.	Other
5.9	17.6	11.8	76.5	11.8	11.8: <ul style="list-style-type: none"> • Time limit is a month in general, can be extended on request • If the users need more than half of the computing power, they have to ask the agreement of the management committee

For a total of 17 complete answers and 1 skipped

3. What do you do regarding accounting, if the job is interrupted because of a server's or service's failure?

Count the job to resource usage limits (as not crashed job)	Don't count the job to resource usage limits	Allow re-running the job apart from limits	re-running the job, but using some smaller system	Other
11.8	76.5	17.6	0	5.9: <ul style="list-style-type: none"> We don't count the job, and if possible the job is re-running by the administrators

For a total of 17 complete answers and 1 skipped

4. What will you do on PETAFL0P systems regarding accounting, if the job is interrupted because of a server's or service's failure?

Count the job to resource usage limits (as not crashed job)	Don't count the job to resource usage limits	Allow re-running the job apart from limits	allow re-running the job, but using some smaller system	Other
0	93.3	26.7	0	6.7 Count some portion of the job, with an assumption that the job will have regular application-level checkpointing.

For a total of 15 complete answers and 3 skipped

5. Do you have your own personnel that take cares of the system administration/maintenance?

Yes	No
89	11

For a total of 17 complete answers and 1 skipped

6. Will you take care of the PETAFL0P system administration/ maintenance in your own?

Yes	No	Other with comment like "additional on-site personnel from vendor"
93	6	1

For a total of 13 complete answers and 5 skipped

Local site allocation policy example (Service Level Agreement), by the 'local site' we understand the independent infrastructure assembled in a given HPC centre here:

- Assumption
 - Local Ressource Management System (LRMS) has agreed to make Resource available to User for a period of one day.
- How is this agreement to be interpreted?
 - Resource might be given to User if needed.
 - LRMS might make Resource available to others when User is not using it.
 - Site might commit to preempt other users as soon as User requests Resource.
- Over-usage
 - If User is allowed to acquire more than agreed Resources, then this may or may not result in User's allocation being reduced later in the given period of time.

LRMS here can correspond to the distributed RMS as well with all the consequences regarding: processing jobs, resources allocation and timeline.

What do we do regarding accounting, if the job is interrupted because of a server's or service's failure. Only 76% of the responders say that they do not count the job resource utilisation to the granted thresholds. It covers the general view of not paying for unsatisfactory level of service. All centres supports the crashed jobs by the rerunning on the same machine or smaller one if applicable. Most of such jobs are put again on top of the pending jobs' queue.

Notwithstanding some of PRACE partner's centres collect and process all resource usage and bill the users (asking for the e.g. scientific report). There should be a trade off worked out on the uniform treatment of the 'lost' computational power caused by the non user-caused failure.

For the forthcoming systems (production and later), there is a common point of interest not to count the "crashed jobs" resource usage (close to 93% responses). Re-running of the job should be possible.

Site job operation examples concerns management the jobs with over-limit requirements. We asked some questions regarding noteworthy subject:

- How do we handle very LARGE jobs?
- How do we handle very LONG jobs?
- What do we do regarding accounting, if the job is interrupted because of a server's or service's failure?

The final deliverable for M24 should also contain the policy for the following problems:

- How usage policies are enforced at the resource and PRACE RI level?
- What strategies must PRACE deploy to ensure usage policy enforcement?
- How are usage policies distributed to the enforcement points?
- How usage policies are made available to PRACE job and data planners?

7.2 Main processes

7. What processes do you have well defined?

Peer review	Grant Evaluation	Resource evaluation	System Administration	Periodic resource utilisation checking - requesting for the usage summary (eg. Yearly)	Revoking the permissions due to the breaking the r-limits	Other
47.1	52.9	88.2	94.1	88.2	47.1	5.9: We have management committee reviews and users committee reviews each 3 months

For a total of 17 complete answers and 1 skipped

8. Are there any additional processes should be defined in the PETAFL0P environment?

Only one answer received (17 skipped):

- Public Relations, Dissemination

9. What are the main parts of your organisation?

Management board	Peer review team	Administrator team	The support and training team	Dissemination team	Other with answers like
76.5	47.1	88.2	70.6	23.5	35.3: <ul style="list-style-type: none"> • Research teams • Networking • Developers, project workers • System and operations team

For a total of 17 complete answers and 1 skipped

In question 7 the outcome shows that currently in the HPC centres several internal processes are well defined (starting from the most common):

- System administration.
- Resource allocation.
- Periodic resource utilisation checking - requesting for the usage summary (e.g. yearly).
- Grant evaluation.
- Peer review.
- Revoking the permissions due to the breaking the r-limits.
- Management committee reviews and users committee reviews.

The deliverable D2.1.1 Report on options for a legal entity points also processes regarding: coordination, procurement, human resources management, legal support, business development, operations, education and collaboration, applications and possibly computer science and other R&D. In detail the PRACE processes will address the final usage model defined in the deliverable D2.3.2 Usage Model document. The usage model identifies three access modes to the facilities for academic researchers once PRACE is established: Preparatory Access, Project Access, and Programme Access. The details of how the systems would operate to be able to deliver these approaches will be covered in final deliverable D2.6.2.

Current peer review practices are collected by the task 2.4 of WP2, and analyzed in the deliverable D2.4.1 Initial Report on the Peer Review Process. The peer review process consists of several sub processes. In general the peer review system for PRACE will ensure a one-stop shop for access to European ITC resources. The process is developed with the guiding principle for resource allocation: scientific quality of the proposal and the need to use the tier-0 system. The HET (High Performance Computing in Europe Taskforce) peer review recommendations are taken into account during designing of the process.

7.3 Main parts of organisation

The above-mentioned processes may to be operated by PRACE staff, as it is proposed so called PRACE Office.

The main parts of current HPC centres include:

- Management Board
- System administration Team
- Support & Training Team
- Peer Review Team

- Dissemination Team

The responders of the survey also mentioned on question 9:

- Development Team
- Networking Team
- Operations Team
- Research Team

Regarding the peer review it is proposed that there will be a PRACE Office to manage of all the aspects of the peer review process, including the proposal submission process, the servicing of committees/panels, and issuing of calls. The early proposition of the peer review process defines the set of bodies to perform certain actions: Technical Assessment Team, Scientific Assessment Team, Panel for Prioritisation of the Proposals, and Founding Decision Team. It is also proposed in the deliverable D2.4.1 that PRACE will separate the funding decisions from the peer review process, e.g. those acting, as peer reviewers will not also be responsible for authorising the funding decision. Similar approach is also proposed by HET and already used by several countries. The precise nature of the committee/panel structure and their responsibilities will be developed in D2.4.2 Peer review process document.

7.4 System administration

System administration is one of the core processes in the HPC centre, and it will be such in future PRACE RI. Some of the survey (few of them) responders stated they outsource the administration of the systems to external companies and they plan to continue this approach for the future petaflops installations. In such cases any recommendations for the system administration process given by the PRACE RI could be considered as a part of an agreement with the external company.

The core operational tools currently used in HPC centres are identified and described in WP4 deliverable D4.1.2 “Report on existing tier-0 Systems Management solutions”. There are strong dependencies between the final set of tools and solutions, which will be chosen for PRACE, and the process of PRACE systems administration.

On the PRACE level the system administration could cover some common tasks as a central service:

- user authentication procedures and mechanisms,
- information provisioning about configuration and availability of resources,
- central point of access to the user support services.

The site (HPC centre) which will provide the computational resources will be expected to meet requirements and recommendations from the PRACE Work Packages, including:

- deploy the common operational procedures for PRACE,
- collect and maintain information on the services they provide,
- deploy the operation and resource allocation policies,
- deploy a set of middleware and site management tools,
- respond to appropriate tickets (user support).

The intermediate outcome from Question 1 of the section on the survey regarding operation on the very large jobs (in terms of the required storage) is that most centres agree for special usage of the space in dedicated LRMS queues (around 76% for both large and long jobs). In some centres, users need to get the additional agreement from the higher body (like Management Committee) for extra resources. 35% cases indicate that such computational grants get higher priority and extra disk space in case of large jobs (only 12% for long jobs).

Only one answer announces, that very long jobs are not supported. We do not formalise the definition, what the 'long' term means.

7.5 Critical delivery time

“Just on time” completion of the job is not considered here, but we remain it for the M24 outcome. Here, we make the remark that there is such type of the jobs with exact time constraints (e.g.. weather forecasting, sea wave modelling, computer aided medical treatment).

7.6 Initial specifications for the Back office services

Following our discussions and the feedback obtained from the survey we are proposing this preliminary list of specifications for the Back Office services:

- Large or long simulations are handled by dedicated batch queues
- For the forthcoming systems (production and later), there is a common point of interest not to count the “crashed jobs” resource usage due to server or service’s failure
- Close to all the centres have their own personnel for administration/maintenance
- As complement of the proposed PRACE Office on each HPC Centre with the following staff:
 - Management Board
 - System administration Team
 - Support & Training Team
 - Peer Review Team
 - Dissemination Team

The outcome of the survey showed need to have a Research Team, a Development Team (both of them may be merged together), an Operation Team and a Networking Team (which may be merged with the System administration Team).

- Once PRACE will be established the final usage model (which identifies three access modes to the facilities for academic researchers) will be précised on D2.3.2.
- Regarding the peer review process, the precise nature of the committee/panel structure and their responsibilities will be developed in D2.4.2 *Peer review process document*.

8 Management and decision support

8.1 Management and decision support

Upgrading or replacing a large HPC system is a major event in the operation of a computer centre. In most cases, it involves a decision taken at the level of a management structure that supervises the computer centre and the allocation of a specific budget for this evolution. Such a decision is taken usually on the basis of:

- a scientific case based on the requirements of the users;
- strategic goals based on scientific objectives, technology opportunities as well as international and national positioning;
- a business case proving that the allocated budget is compatible with the performance target.

It is important to consider different elements and their sequence and interaction when managing a computing centre setup or upgrade. We briefly consider this in section 8.2.

As those operations are supposedly repeatable or cyclic, it is equally important to organize monitoring and feedback to continuously improve the process on the long term. We insist on those issues on section 8.3.

Below we will present the results for the questionnaire's questions relatives to Management and decision support issues.

1. How often do you perform procurement actions?

Once a year	Monthly	at already predefined times	When needed	Other with answers like
11	0	5	80	5: <ul style="list-style-type: none"> • Typically we would procure a new system every two to three years, with upgrade options in the contracts to avoid the costly business of the procurement process. • HPC systems once a year

For a total of 17 complete answers and 1 skipped

2. Do you coordinate procurement of new systems among sites?

No	Yes, on university side level	Yes, on regional level	Yes, on national level	Who is responsible for performing the actions?
47.1	11.8	0	41.2	<ul style="list-style-type: none"> • I coordinate procurement only on CERFACS site • Sometimes during project, coordinator of the project is usually responsible • Handled by research councils • UNINETT Sigma and HPC/data centres • Gauss Centre for Supercomputing • Management committee for budget and planning, computing

				centre for technical specifications, analysis, and technical choice. <ul style="list-style-type: none"> • CSC • GENCI
--	--	--	--	---

For a total of 17 complete answers and 1 skipped

3. To whom do you need to report on?

- 100% for Usage of your computing infrastructure with answers like:
 - CERFACS' director
 - Advisory Board and management board of parent institution (ETHZ).
 - SNIC (national centre)
 - Ministry of Science and Higher Education
 - Ministry
 - Management board
 - Research councils
 - Research Council of Norway
 - Steering Committee and Commission for Informatics
 - CNRS direction
 - Polish Ministry of Science and Higher Education

For a total of 17 complete answers and 1 skipped

4. How often do you need to report on: usage of your computing infrastructure (i.e. system utilisation) or scientific quality of computation results, other?

Monthly	Yearly	Not at all	On demand	Other
17.6	70.6	0	41.2	20: <ul style="list-style-type: none"> • Quarterly usage statistics are provided in addition to annual report. • Every 6 months • 2x per year (I had to pick one of the answers here, to avoid error) • Each 3 months for usage of computing infrastructure and each year for scientific results

For a total of 17 complete answers and 1 skipped

5. Are you given the fixed report template document?

Yes	No
29.4	70.6

For a total of 17 complete answers and 1 skipped.

8.2 Planning and procurement management

The deliverable D2.3.1 and the following deliverables that will be produced by WP7-task6 focus on the procurement process for an HPC system. D2.3.1 describes, on the basis of what is usually done in most European HPC centres, a generic process divided in 8 stages while WP7 deliverables address the practical implementation of the process for the procurement of an HPC system of the future RI.

Overall process

The procurement process is indeed a major component of the overall planning of a large supercomputer installation or upgrade. However, this planning must take into account other important elements including:

- facility construction or upgrade;
- installation or upgrade of other IT equipments (for example: long term storage system);
- preparation of the production including:
 - adjusting the system administration and application support teams sizes and skills to the new supercomputer,
 - integration in data centre IT infrastructure
 - adaption of data centre policies and deployment of the system software for their enforcement
 - strategy for switching from one supercomputer to the other (overlap, ...),
 - migration of user data (if needed),
 - application software porting planning (migration/optimisation of applications towards the new system).

Since these elements are closely related to each other they need to be planned and controlled by a single project team to ensure overall consistency and efficiency. This is also necessary since all these elements contribute to the TCO of the computer centre.

In Europe, the ways these processes are organized differ a lot from one country to the other. In some countries (for example UK), most of the work is subcontracted to a single supplier (usually a consortium) under a single contract, while in other countries (for example France and Germany) the computer centre team directly manages and coordinate the various tasks.

Facility construction or upgrade

Petascale supercomputers have very demanding installation requirements, particularly in terms of power consumption (several MW) and cooling (air and/or water). This means that the technical facilities get more and more expensive and complex.

In addition to cost, timing is an important issue: it would be desirable to start building the technical facilities only after the choice of the supercomputer that will be installed but that would mean a long time (6 to 18 months³) between the selection of the supercomputer and its installation. Since this is usually not possible, the way to deal with this issue is in general:

³ Depending whether only limited evolution of existing facilities is necessary or if new large equipments must be installed (chillers, UPS, water loops, ...).

- to perform a market survey in order to estimate the installation requirements before starting the procurement
- to start the construction of the technical facilities according to these estimates
- to describe the already known technical facilities features in the technical requirements of the procurement and to require that the vendor proposes systems comply with these facilities.

Of course the distribution of electricity, air and if needed water to the supercomputer can't be prepared before the choice of the supercomputer to install. However, the corresponding delays are shorter since lightweight procurement processes can then be used (smaller amount of money involved) and because the construction is simpler.

Installation or upgrade of other IT equipments.

In order to keep the computer centre architecture balanced, it is necessary to ensure a good balance between the storage system of the computer centre (of greater capacity but slower bandwidth than the storage directly attached to the supercomputer), the internal network capacity and the performances of the supercomputer.

Typically, a recommended capacity for the storage system is 10 times the capacity of the supercomputer (with a similar ratio, in the opposite direction, for the bandwidth).

Preparation of the production

In order to ensure proper production preparation, it is desirable to put in place a user group that will meet on a regular basis. The meetings usually start before the production and continue until the major applications have been ported and optimized.

8.3 Monitoring and reporting

Once the decision has been implemented (supercomputer installed and running), the computer centre reports on a regular basis, in particular to the management structure that supervises it. This reporting is useful both for measuring how well the new HPC system is performing and for preparing further evolutions.

For the majority of European HPC centres, there are three main aspects to reporting:

- 1) Quality of service.
The most important parameter in this context is the availability of the system, which measures the stability system and the quality of system management. The assessment must distinguish between unplanned interrupts and planned interrupts. User satisfaction surveys are an additional important method to measure QoS.
- 2) Levels of utilisation of the system.
The frequency of reporting varies across the centres, ranging from monthly to yearly. In a small number of cases, the centres reported on demand from the management structures. Includes parameters like job size distributions, job durations, distribution by application field and user institution.

- 3) On the scientific quality of the results of the computational experiments.
Based on typical scientific indicators like bibliometric parameters, scientific awards and prizes, etc. The frequency of reporting varies across the centres, ranging from monthly to yearly. In a small number of cases, the centres reported on demand from the management structures. Typically these indicators refer to usage of the HPC service one to three years ago.

Apart from a small number of cases, it is usual for the centres to decide the format of the reports themselves – they are not provided with a fixed report template.

8.4 Initial specifications for Management and decision support

Following our discussions and the feedback obtained from the survey we are proposing this preliminary list of specifications for the Management and decision support:

- In order to ensure overall consistency and efficiency during a procurement, all the items listed on the procurement planning must be managed by a single project team
- Other IT equipments (storage, network, ...) have to be taken into account in order to provide a balanced computing environment
- In order to ensure proper production preparation, it is desirable to put in place a user group that will meet on a regular basis
- Regular Monitoring and reporting must be performed with the following 3 targets: Quality of Service, Level of utilisation of tier-0 systems and scientific quality of the results of the computational experiments

9 Conclusion and next steps

After having conducted an intensive survey among European HPC tier-1 and tier-2 centres (with 18 answers) we've have been able to have a detailed overview of yet existing services gathered on 6 main sections: Front Office, User Support/Training, Security, Ecosystem Insertion and Back Office and Management/decision support.

We've also gathered projection of services on future tier-0 systems, which allows us with the output of the other connected deliverables to propose some preliminaries initial specifications for each listed service.

The following D2.6.2, targeted for month 24, will give a complete but not final specification for the PRACE operational model.

It will take as an input the initial description done in 2.6.1 and through work/results of deliverables of PRACE's technical work packages (especially WP4, WP5, WP7 and WP8) and other upcoming deliverables from WP2 and WP3 will define the "best set" of services delivered by the 3 to 5 HPC infrastructures which will be deployed starting 2010.

10 Annex

10.1 Description of the survey

In order to have a detailed overview of existing services implemented by European HPC centres an extensive electronic survey has been conducted during a period of 3 weeks. The invitation to participate to the survey has been sent to all the PRACE partners involved as tier-1 on each their own country. Most of these centres are also involved on other European HPC projects like DEISA or HPC Europa with an existing structure of services offered to users.

This survey has also been sent to some tier-2 like CERFACS, which have the particularity to address HPC services to industry.

The questionnaire has been published using the SurveyMonkey tool (<http://www.surveymonkey.com>) on 17th November 2008. It has been closed on Friday December 5th 2008.

10.2 Survey answers

Almost 18 answers have been collected after this period.

Name	Site	Country	Type
Sergi Girona	BSC	Spain	tier-1 and tier-2
Bartosz Pliszka	TASK	Poland	tier-1
Francis Daumas	CINES	France	tier-1
Victor Alessandrini	IDRIS	France	tier-1
Jussi Heikonen	CSC	Finland	tier-1
Christine Menaché	CCRT	France	tier-1
Pedro Alberto	Universidade de Coimbra	Portugal	tier-1
Martin Polak	GUP	Poland	tier-1
Mirosław Kupczyk	PSNC	Poland	tier-1
Matthias Brehm	LRZ	Germany	tier-1
Jacko Koster	UNINETT Sigma	Norway	tier-1
Stephen Booth	EPCC	UK	tier-1
Sergio Bernardi	CINECA	Italy	tier-1
Maciej Filocha	ICM UW	Poland	tier-1
Agnieszka Kwiecien	WCSS	Poland	tier-1
Erwin Laure	KTH/PDC	Sweden	tier-1
Neil Stringfellow	ETHZ-CSCS	Switzerland	tier-1
Nicolas Monnier	CERFACS	France	tier-2

10.3 Survey details and results

Introduction

Question 1: Information about the respondent

First Name

Surname

Affiliation

HPC & Related Services

1. Is an access to HPC resources uniform per infrastructure or site?

Per infrastructure

Per site

The access is specific

The answers received showed the following results

- 47% per infrastructure
- 47% per site
- 5.3 % specific

For a total of 18 complete answers

2. What type of access to resources is available to users?

Command line

Web

GUI

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 100% for the Command line mode
- 26.3% for the WEB mode
- 15.3% for the GUI mode
- 15.3 % for Other with 2 answers with Globus/Unicore middleware

For a total of 18 complete answers

3. Is the access restricted in some way?

Registered subnet

Registered machine IP

Allowed from any machine

Other (please specify)

The answers received showed the following results

- 15% for an access through Registered subnet
- 26% for an access through Registered machine IP
- 58% for an access allowed from any machine

For a total of 18 complete answers

4. Is it possible for a user to submit, monitor and manage jobs in a uniform way common user environment, or any integrated environment? {Current state}

Yes, per infrastructure

Yes, per site

Only submit per infrastructure

But not monitoring and managing

Other (please specify)

The answers received showed the following results

- 47% per infrastructure
- 47% per site
- 15.8% only submit per infrastructure
- 5.3% but not monitoring and managing

For a total of 18 complete answers

5. Should it be possible for a user to submit, monitor and manage jobs in a uniform way - common user environment, or any integrated environment? {For future Petaflop system} state}

Yes, per infrastructure

Yes, per site

Only submit per infrastructure

But not monitoring and managing

Other (please specify)

The answers received showed the following results

- 60% per infrastructure
- 30% per site
- 10% Only submit per infrastructure

Some remarks about the wish to have a PRACE uniform submit/monitor way

For a total of 18 complete answers

6. What kinds of jobs are available to users now?

Batch

Interactive (online job control, eg visualisation)

Local

Remote

As answers were not exclusive, The answers received showed the following results

- 94.4 % in batch mode
- 68.4% in interactive mode
- 26.2% in local mode
- 36% in remote mode

For a total of 18 complete answers

7. What kind of jobs should be available to users? {in future petaflop system}

Batch

Interactive (online job control, eg visualisation)

Local

Remote

As answers were not exclusive, the answers received showed the following results

- 100 % in batch mode
- 75% in interactive mode
- 37.5% in local mode
- 62.5% in remote mode

For a total of 16 complete answers and 2 skipped answers

8. Is the job checkpointing available to the users now?

Yes, systemise

Yes, user wise

No

No, because we rely on UPS, generators, etc...

As answers were not exclusive, the answers received showed the following results

- 5.6% for system wise
- 47.4% for user wise
- 42.1% for no checkpointing
- 5.8% rely on their UPS so no checkpointing

For a total of 18 complete answers

9. Is the job checkpointing necessary in the Petaflop installation?

Yes, systemise

Yes, user wise

No

No, because we rely on UPS, generators, etc...

As answers were not exclusive, the answers received showed the following results

- 35.3% for system wise
- 52.9% for user wise
- 23.5% for no checkpointing
- 11.8% rely on their UPS so no checkpointing

For a total of 17 complete answers and 1 skipped answer

10. Is the pending job rerouting available to the users (can users change the queue manually)?

From one queue to another on the same machine

From one system to another

No

The answers received showed the following results

- 26.3% From one queue to another on the same machine
- 0% From one system to another
- 73.7 For the No answer

For a total of 18 complete answers

11. Should the pending job rerouting be available to the users (Would users change the queue manually) in the petaflop machine?

From one queue to another on the same machine

From one system to another

No

The answers received showed the following results

- 46.7% From one queue to another on the same machine
- 26.7% From one system to another
- 46.7% For the No answer

For a total of 15 complete answers and 3 skipped answers

12. Is the resource reservation available to the users?

Yes

No

The answers received showed the following results

- 52.6% for the Yes answer
- 47.4% for the No

For a total of 18 complete answers

13. Should the resource reservation be available to the users in petaflop systems?

Yes

No

The answers received showed the following results

- 88.2% for the Yes answer
- 11.8% for the No

For a total of 17 complete answers and 1 skipped answer

14. Are workflow or co-allocation capabilities available to the users?

Yes

No

The answers received showed the following results

- 47.4% for the Yes answer
- 52.6% for the No

For a total of 18 complete answers

15. Should workflow or co-allocation capabilities be available to users in petaflop systems?

Yes

No

The answers received showed the following results

- 87.5% for the Yes answer
- 12.5% for the No

For a total of 16 complete answers and 2 answers skipped

16. Is a set of applications homogeneous between sites/servers?

Yes

No

Some

The answers received showed the following results

- 26.3% for the Yes answer
- 5.3% for the No
- 78.9% for the Some answers

For a total of 18 complete answers

17. Should a set of applications be homogeneous between sites/servers in petaflop environment?

Yes

No

Some

The answers received showed the following results

- 17.6% for the Yes answer
- 17.6% for the No
- 64.7% for the Some answers

For a total of 17 complete answers and 1 answer skipped

18. How the software licenses are managed?

Locally

Shared between sites

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 84.2% for the Locally answer
- 26.3% for the Shared
- 13.4% for the Other

For a total of 18 complete answers

19. How the software licenses should be managed?

Locally

Shared between sites

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 71.4% for the Locally answer
- 50% for the Shared answer
- 32% for the Other

For a total of 14 complete answers and 4 skipped answers

20. Do you use the federated accounting data for purposes of?

Billing

Keeping the resources thresholds

For watching (for fun)

No federated accounting

Other (please specify)

As all answers were not exclusive, the answers received showed the following results

- 42.1% for the Billing answer
- 73.7% for the answer “Keeping the resources thresholds
- 15.8% for “For watching (for fun)”
- 10.5% for the “No federated accounting” answer
- 5.5% for Other answer

For a total of 18 complete answers

21. Will you use the federated accounting data from petaflop systems for purposes of?

Billing

Keeping the resources thresholds

For watching (for fun)

No federated accounting

Other (please specify)

As all answers were not exclusive, the answers received showed the following results

- 71.4% for the Billing answer
- 85.7% for the answer “Keeping the resources thresholds
- 14.3% for “For watching (for fun)”

- 0% for the “No federated accounting” answer
- 28.6% for Other answer

For a total of 14 complete answers and 4 skipped answers

22. Are there dedicated service/services for users/groups for data upload/download and management?

Across infrastructure

Across sites

Per server

As all answers were not exclusive, the answers received showed the following results

- 42.1% for the Across infrastructure
- 36.8% for the Across sites answer
- 36.8% for the Per Server answer

For a total of 18 complete answers

23. Should dedicated service/services for users/groups for data upload/download and management be available on petaflop environment?

Across infrastructure

Across sites

Per server

As all answers were not exclusive, the answers received showed the following results

- 47.1% for the Across infrastructure
- 58.8% for the Across sites answer
- 23.5% for the Per Server answer

For a total of 17 complete answers and 1 skipped answer

24. Do you provide any specialized services for visualisation of the computations?

Yes

No

Other (please specify)

The answers received showed the following results

- 63.2% for the Yes
- 26.8% for the No

For a total of 18 complete answers

25. Would the provision of any specialized services for visualisation of the PETAFLOP computations be necessary?

Yes

No

Other (please specify)

The answers received showed the following results

- 75% for the Yes
- 25% for the No

For a total of 12 complete answers and 6 skipped answers

26. Are users allowed to make a backup of their data on storage systems?

If yes what are the parameters of the service? (Size; time frame)

Average data size

Time (eg. during the project duration, after completion X months,...)

No (rewrite 'no')

As answers were not exclusive, the answers received showed the following results

- 68.4% for the Average Data size with a quite vast distribution (from 5 GB to unlimited!)
- 63.2% for the Time answer (from one year to infinite through HSM)
- 31.6% for the No answer

For a total of 18 complete answers

27. Should the users be allowed to make a backup of their data on storage systems in their own (in PETAFLOP environment)?

If yes what are the parameters of the service? (Size; time frame)

Average data size

Time (eg. during the project duration, after completion X months, ...)

No (rewrite 'no')

As answers were not exclusive, the answers received showed the following results

- 64.3% for the Average Data size with a quite vast distribution (from 500 GB to unlimited!)
- 64.3% for the Time answer (from 3 months to project life, without HSM)
- 28.6% for the No answer

For a total of 14 complete answers and 4 skipped answers

28. Is there an infrastructure dedicated for testing and development available to users? (yes; no) If yes what kind?

Yes - dedicated short queues on production systems

Yes - exclusive node(s) or partition on the production systems

Yes - other compatible machine

Yes - interactive testing

No

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 68.4% for “Yes - dedicated short queues on production systems”
- 42.1% for “Yes - exclusive node(s) or partition on the production systems”
- 15.8% for the “Yes - other compatible machine”
- 57.9% for the Yes – Interactive testing
- 21.1% for the No

For a total of 18 complete answers

29. Is the special infrastructure for testing and development necessary in PETAFL0P environment? (yes; no) If yes what kind?

Yes - dedicated short queues on production systems

Yes - exclusive node(s) or partition on the production systems

Yes - other compatible machine

Yes - interactive testing

No

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 75% for “Yes - dedicated short queues on production systems”
- 56.3% for “Yes - exclusive node(s) or partition on the production systems”
- 18.8% for the “Yes - other compatible machine”
- 43.8% for the Yes – Interactive testing
- 6.3% for the No

For a total of 16 complete answers and 2 skipped answers

30. Is there a shared workspace available for users?

Yes - it is shared between many systems in a single site

Yes - it is shared among multiple sites

No

The answers received showed the following results

- 83.5% for Yes - it is shared between many systems in a single site
- 5.3% for Yes - it is shared among multiple sites
- 10.5% for the No answer

For a total of 18 complete answers.

31. Should a shared workspace be available for users in the PETAFL0P environment?

Yes - it is shared between many systems in a single site

Yes - it is shared among multiple sites

No

As answers were not exclusive, the answers received showed the following results

- 75% for Yes - it is shared between many systems in a single site
- 50% for Yes - it is shared among multiple sites
- 0% for the No answer

For a total of 16 complete answers and 2 skipped answers

Ecosystem integration

Questions regarding Network Connectivity between PRACE's tier-0 sites and national tier-1 (or regional tier2) or other actors as described on 2.5.1

1. What kind of interoperation do you perform with other sites?

Common services
 Common procedures
 User exchange
 Job Exchange
 Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 63.2% for Common services
- 47.4% for Common procedures
- 42.1% for User exchange
- 21.1% for Job Exchange
- 10% for other (in fact a no answer)

For a total of total of 18 complete answers

2. What kind of interoperation will you intend to perform with other sites in the petaflop environment?

Common services
 Common procedures
 User exchange
 Job Exchange
 Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 71.4% for Common services
- 78.6% for Common procedures
- 57.1% for User exchange
- 35.7% for Job Exchange
- 0% for other (in fact a no answer)

For a total of total of 14 complete answers and 4 skipped answers

3. Do you interoperate with some grid infrastructure? (Yes, no)

Yes, which one?

Purpose of interoperation

No (rewrite 'no')

As answers were not exclusive, the answers received showed the following results

- 71.4% for Yes with Synfiniway, EGEE framework, DEISA, globus, unicore, Chemomomentum, Spanish Supercomputing Grid, and Nordic Data Grid Facility
- 78.6% for Purpose of operations with answers like
 - access from Airbus to a BlueGene system
 - Offering resources to larger user communities, providing local user communities with access to other resources
 - Ressources utilisation
 - Common research projects
 - To execute jobs on HPC machines otherwise not available to Italian HPC users
 - Supports complex workflow or portal interface to standard packages
 - Distributed Nordic tier-1
 - Resource exchange, benchmarking; LHC research
 - For experimental purpose
- 21.1% for the No answer

For a total of 18 complete answers

4. Can you see the advantage of interoperation petaflop systems with some grid infrastructure? Technically please!!

Yes - which grid infrastructure?

Purpose of interoperation

No (rewrite 'no')

As answers were not exclusive, the answers received showed the following results

- 66.7% for Yes with answers like Deisa, globus, Chemomentum, D-Grid and EGI
- 60% for the Purpose of Operations with answers like
 - Workflows (locally pre-process, tier-0 compute, locally or tier-0 post-process)
 - Remote workflow, access to chemical (toxicity) databases
 - To implement the PRACE vision (pyramid of tier-0 and tier-1)
 - Requirement of some (but not all) user groups
 - Easy migration of jobs and data between tier-0 and tier-1 systems
 - Common user environment
 - Submitting jobs for farming applications
 - AAA, data movers, current DEISA users are the most potential petaflop users

For a total of 15 complete answers and 3 skipped answers

User Support

Questions about user support

1. What kinds of interfaces are available to users to request for support?

Web

Email

Phone

Service Desk

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 61.1% for Web
- 100% for Email
- 83.3% for Phone
- 44.4% for Service Desk

For a total of 18 complete answers

2. What is the workflow and responsibilities to solve the users' problems? (i.e. requests are forwarded from the support team for the infrastructure to the site administrators, then to a system or application administrator; to other users; other)

Central point of requests

Direct requests to the administrator of (system, application, network, security,...)

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 77.8% for Central point of requests
- 38.9% for Direct requests to the administrator of (system, application, network, security,...)

For a total of 18 complete answers

3. Are the user support processes certified or implemented in accordance with standards? (i.e. ISO 20000, or other) If yes, which one?

Yes, which one

No

The answers received showed the following results

- 27.8% for the Yes answers with ISO9001 or ISO27000 or ITIL
- 38.9% No

For a total of 18 complete answers

4. Do you think, that the user support should be implemented in accordance with standards in petaclass environment? (i.e. ISO 20000, or other) If yes, which one?

Yes, which one

No

The answers received showed the following results

- 33.3% for the Yes answers with ISO9001 or ISO27000 or ITIL
- 66.7% No

For a total of 12 complete answers and 6 skipped

5. What user support services are available now?

Help with preparation of data for computation

Help with post-processing, i.e. visualisation

Parallelisation of application code (from sequential handed)

Help with tuning of parallelized application

Help with tuning of any application code

Help with algorithms

Help with numerics

Help with porting applications

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 66.7% for Help with preparation of data for computation
- 61.1% for Help with post-processing, i.e. visualisation
- 61.1% for Parallelisation of application code (from sequential handed)
- 94.4% for Help with tuning of parallelized application
- 77.8% for Help with tuning of any application code
- 33.3% for Help with algorithms
- 33.3% for Help with numerics

- 94.4% for Help with porting applications
- 6% for Other with science support (e.g. recommending software packages)

For a total of 18 complete answers

6. What user support services should be available in the future petaclass environment?

Help with preparation of data for computation

Help with post-processing, i.e. visualisation

Parallelisation of application code (from sequential handed)

Help with tuning of parallelized application

Help with tuning of any application code

Help with algorithms

Help with numerics

Help with porting applications

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 86.7% for Help with preparation of data for computation
- 80% for Help with post-processing, i.e. visualisation
- 80% for Parallelisation of application code (from sequential handed)
- 100% for Help with tuning of parallelized application
- 66.7% for Help with tuning of any application code
- 80% for Help with algorithms
- 73.3 for Help with numerics
- 93.3% for Help with porting applications

For a total of 15 complete answers and 3 skipped

Training

Questions relative to training activities.

1. What kind of training is available to users?

Group

Personal

Periodical

On demand/request

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 76.5% for Group
- 23.5% for Personal
- 41.2% for Periodical
- 52.9% for On Demand/request
- 6% for www training material

For a total of 17 complete answers and 1 skipped

2. What kind of training should be available to users in the petaclass environment?

Group

Personal

Periodical

On demand/request

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 92.9% for Group
- 14.3% for Personal
- 64.3% for Periodical
- 64.3% for On Demand/request
- 6% for www training material

For a total of 14 complete answers and 4 skipped

3. Are the courses available only to the site's users or are open for others?

Open – Free

Open – Higher fee

Only site registered people

Other conditions for non-registered people:

As answers were not exclusive, the answers received showed the following results

- 92.9% for Open-Free answer
- 14.3% for Open-Higher fee answer
- 11.8% for Only site registered people

For a total of 17 complete answers and 1 skipped

4. Should the courses be available only to the site's users or left open for others in the future?

Open – Free

Open – Higher fee

Only site registered people

Other conditions for non-registered people:

As answers were not exclusive, the answers received showed the following results

- 86.7% for Open-Free answer
- 20% for Open-Higher fee answer
- 6.7% for Only site registered people

For a total of 15 complete answers and 3 skipped

5. Are your hardware/system/application vendors performing training for the users?

Yes

No

Other (please specify)

The answers received showed the following results

- 82.4% for Yes
- 23.5% for No
- 23.5% for Other with answers like
 - One or two initial courses are provided in the first year of operation
 - Yes, but the effort spent on this is limited ,
 - Depends on needs!
 - Additional training by vendors

For a total of 17 complete answers and 1 skipped

6. Should your hardware/system/application vendors perform training for the users in the future?

Yes

No

Other (please specify)

The answers received showed the following results

- 85.7% for Yes
- 14.3% for No

For a total of 14 complete answers and 4 skipped

7. Do you have a training program for your personnel?

Yes

No

The answers received showed the following results

- 47.1% for Yes
- 52.9% for No

For a total of 17 complete answers and 1 skipped

8. Are you willing to have a training program for your personnel in the future petaflop environment?

Yes

No

The answers received showed the following results

- 93.3% for Yes
- 6.7% for No

For a total of 15 complete answers and 3 skipped

9. Provide a list of courses that are available at your site (or URL to the list available on-line)

Titles or URL

12 complete answers with:

- We don't have any list on our site
- We have an annual MPI/OpenMP course given by Ralf Rabenseifner (HLRS), and we are currently preparing our full course program of other material.
- <http://www.pdc.kth.se/events/events2008/view>
- Courses are prepared on demand
- <http://www.icm.edu.pl/kdm/Szkolenia>
- <http://www.cineca.it/scuoleestive/>
- <http://www.hpcx.ac.uk/support/training/index.html>
<http://www.hector.ac.uk/support/cse/schedule/>
- <http://www.lrz-muenchen.de/services/compute/courses/>
- <http://szkolenia.man.poznan.pl>
- The list is available on the CCRT web site

Dissemination

Questions relatives to dissemination activities (seminar, web, regular newsletters, forums, ...)

1. What kind of dissemination actions do you perform regarding your infrastructure or site and its usage? Provide a list of activities.

Conference papers

Stands / exhibitions

International/local events about computing

User forums

Journals

Printed newsletters

Advertising

Other (please specify)

The answers received showed the following results

- 52.9% for Conference papers
- 52.9% for Stands/Exhibitions
- 76.5% for International/local events about computing
- 58.8% for User forums
- 29.4% for Journals
- 35.3% for Printed Newsletters
- 5.9% for Advertising
- 12% for Other (no dissemination actions and Electronic Newsletters)

For a total of 17 complete answers and 1 skipped

2. Do you have dedicated personnel to perform dissemination and outreach?

Yes

No

Other (please specify)

The answers received showed the following results

- 58.8% for the Yes answer
- 41.2% for No

For a total of 17 complete answers and 1 skipped

Data security

1. Is there a uniform user authentication and authorisation service available?

Across the infrastructure

Across the site

No

No - distinct services for different types of data (e.g. astrophysical vs. automotive)

As answers were not exclusive, the answers received showed the following results

- 35.3% for the “Across the Infrastructure” answer
- 58.8% for the “Across the site” answer
- 11.8% for the No
- 5.9% for the Other

For a total of 17 complete answers and 1 skipped

**2. Do you allow your users to submit/run their own programs?
If yes, are the applications verified on the level of code, scanned for malicious content etc?**

Yes

If yes, are the applications verified on the level of code, scanned for malicious content, etc?

No

As answers were not exclusive, the answers received showed the following results

- 94.1% for the Yes
- 0%
- 23.5% for the No

For a total of 17 complete answers and 1 skipped

3. Do you have user activity-tracking software on your systems?

Yes, always on

Yes, on demand is on

No

The answers received showed the following results

- 29.4% for the “Yes, always on” answer
- 35.3% for the “Yes, on demand is on”
- 35.3% for the No

For a total of 17 complete answers and 1 skipped

4. Do you provide your users with an SLA regarding the uptime, confidentiality, data consistency etc?

The uptime
 Confidentiality
 Data consistency
 Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 35.3% for the Uptime answer
- 29.4% for the Confidentiality answer
- 29.4% for the Data consistency answer
- 52.9% for the Other answer with:
 - No SLA
 - Only requested from commercial users
 - Best effort
 - Agreement for service downtime twice a month: 1.&3.Thursday.
 - Don't know of any

For a total of 17 complete answers and 1 skipped

5. Are your users responsible for any possible damage caused by their behavior, i.e.: security breach because of trivial password or running malicious software, or giving the credentials to someone else? (yes; no) If yes, what are the possible consequences?

Yes

No

If yes, what are the possible consequences?

The answers received showed the following results

- 82.4% for the Yes with answers like:
 - Giving credentials to someone else results in a life ban from using the Centre's facilities (reminding people of this fact acts as a good deterrent).
 - Block users; there might of course also be legal consequences.
 - Closing an account
 - Consequences are those listed by the Italian laws. From administrative computer crime to criminal computer crime.
 - Loss of access
 - Closing of account, termination of access
 - No further access to the system
 - Revocation of access
- 5.9% for No

For a total of 17 complete answers and 1 skipped

6. What procedures do you implement to prevent and limit undesired user behavior?

On-line training for users about basic security issues

An official "I am fully responsible for what I am doing". Document that anyone willing to use the resources in your centre have to sign (bylaw).

Other (please specify)

The answers received showed the following results

- 17.6% for the “On-line training for users about basic security issues” answer
- 82.4% for the second answer
- 29.4% for the Other with answers like:
 - Activity scans
 - None really
 - Acceptance of the Rules and Regulation with application
 - Guide lines on www and application form

For a total of 17 complete answers and 1 skipped

Site Security

1. Do you have your own data centre room?

Yes

No - we rent the space from a third party company

The answer was 100% for the Yes answer, for a total of 17 complete answers and 1 skipped

2. Is your data centre located in a dedicated building or is it a shared-purpose building (offices, other companies etc)

Dedicated building

Shared-purpose building

The answers received showed the following results

- 35.3% for Dedicated Building
- 64.7% for the Shared purpose building

For a total of 17 complete answers and 1 skipped

3. What are the physical access restrictions, physical means for preventing unlawful access?

The 17 complete answers (1 skipped) were the following

- Badge system
- Standard office security - locked doors, electronic key access. The machine room is protected by strong security doors.
- Door access control, movement sensors
- Token authorisation system
- Electronic card + biometrics
- Guards, special locks, CCTV, alarms,
- Electronic card controlled access to several building rooms
- Details restricted
- Key card / pin code access

- Access keys, doors, video, personnel
- This is a protected site with physical access restrictions
- Security guards that control the access to the building, plus TV monitoring, alarms, etc
- Only authorized personnel, proximity cards, biometric access system

4. What is your policy of granting the access to the data centre room to non-staff persons (e.g. visitors, vendors, facility workers, cleaning personnel)?

The 17 complete answers (1 skipped) were the following

- Badge access restrictions
- Access to the machine room for non-authorized personnel (including staff without machine-room access) requires the machine room doors to be opened by the facility management personnel.
- Decided on demand
- Access only with the personnel
- Passport ID is required
- On request policy, implemented by means of electronic card controlled access
- Case by case. Visitors and one-time workers are always accompanied by staff
- Visitors, vendors, facility workers only accompanied by staff, cleaning personnel has their own set of keys
- We have limited access to the data centre room through badge's reader

5. How the verification procedure of the new users allowed using the systems in your data centre look like?

The user presents their private ID.

There is/are person that contacts all the candidates.

The user presents the certified credentials by their home institution (X509-like solution).

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 58.8% for the user presents their private ID.
- 35.3% for the “There is/are person that contacts all the candidates” answer
- 5.9% for the “The user presents the certified credentials by their home institution (X509-like solution)” answer
- 29.4% for Other with answers like
 - Online registration with approval devolved to project management
 - ssh public/private keys
 - Application form
 - Each organisation has representatives who validate the new accounts

For a total of 17 complete answers and 1 skipped

Operation

1. How do you handle very LARGE jobs?

Do the users need to consult the site's Council of Users (higher body) every time before running the job?

Do the users need to consult the site's administrators every time before running the job?

Are there grants with special privileges i.e. additional disk space, or better Priority? Dedicated queues.

Very large jobs are not supported here.

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 5.9% for the first answer
- 17.6% for the before running the job answer
- 35.3% for the "Do the users need to consult the site's administrators every time before running the job?" answer
- 76.5% for the "Are there grants with special privileges i.e. additional disk space, or better Priority?" answer
- 11.8% for the Dedicated queues.
- 29.4% for the Other with answers like
 - Anyone can run a large job without special permission. Part of the scheduling policy is to weight jobs according to job size in addition to other priority factors.
 - Not clear what very large is - our users can use up to the full capacity of a certain machine
 - Queue structures usually give priority to jobs (e.g. large jobs) that cannot be run at smaller facilities.
 - Prioritisation of large jobs
 - If the users need more than half of the computing power, they have to ask the agreement of the management committee

For a total of 17 complete answers and 1 skipped

2. How do you handle very LONG jobs?

Do the users need to consult the site's Council of Users (higher body) every time before running the job?

Do the users need to consult the site's administrators every time before running the job?

Are there grants with special privileges i.e. additional disk space, or better Priority? Dedicated queues.

Very long jobs are not supported here.

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 5.9% for the first answer
- 17.6% for the before running the job answer
- 11.8% for the "Do the users need to consult the site's administrators every time before running the job?" answer
- 76.5% for the "Are there grants with special privileges i.e. additional disk space, or better Priority?" answer
- 11.8% for the Dedicated queues.

- 11.8% for the Other with answers like
 - Time limit is a month in general, can be extended on request
 - If the users need more than half of the computing power, they have to ask the agreement of the management committee

For a total of 17 complete answers and 1 skipped

3. What do you do regarding accounting, if the job is interrupted because of a server's or service's failure?

Count the job to resource usage limits (as not crashed job)

Don't count the job to resource usage limits

Allow re-running the job apart from limits

Allow re-running the job, but using some smaller system

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 11.8% for the "Count the job to resource usage limits (as not crashed job)" answer
- 76.5% for the "Don't count the job to resource usage limits" answer
- 17.6% for the "Allow re-running the job apart from limits" answer
- 0% for the "Allow re-running the job, but using some smaller system" answer
- 5.9% for the Other with answers like:
 - We don't count the job, and if possible the job is re-running by the administrators

For a total of 17 complete answers and 1 skipped

4. What will you do on PETAFL0P systems regarding accounting, if the job is interrupted because of a server's or service's failure?

Count the job to resource usage limits (as not crashed job)

Don't count the job to resource usage limits

Allow re-running the job apart from limits

Allow re-running the job, but using some smaller system

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 0% for the "Count the job to resource usage limits (as not crashed job)" answer
- 93.3% for the "Don't count the job to resource usage limits" answer
- 26.7% for the "Allow re-running the job apart from limits" answer
- 0% for the "Allow re-running the job, but using some smaller system" answer
- 6.7% for the Other with answers like:
 - Count some portion of the job, with an assumption that the job will have regular application-level checkpointing.

For a total of 15 complete answers and 3 skipped

5. Do you have your own personnel that take cares of the system administration/maintenance?

Yes

No – it is done by an external company

The answers received showed the following results

- 89% for the Yes
- 11% for the No

For a total of 17 complete answers and 1 skipped

6. Will you take care of the PETAFL0P system administration/maintenance in your own?

Yes

No – the external company will do this

Other (please specify)

The answers received showed the following results

- 93% for the Yes
- 6% for the No
- 1% for Other with comment like “additional on-site personnel from vendor “

For a total of 13 complete answers and 5 skipped

Back office

1. What processes do you have well defined?

Peer review

Grant evaluation

Resource evaluation

System administration

Periodic resource utilisation checking - requesting for the usage summary (eg. yearly).

Revoking the permissions due to the breaking the r-limits.

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 47.1% for the Peer review answer
- 52.9% for the Grant Evaluation
- 88.2% for the Resource evaluation
- 94.1% for the System Administration
- 88.2% for the Periodic resource utilisation checking - requesting for the usage summary (e.g. Yearly) answer
- 47.1% for the Revoking the permissions due to the breaking the r-limits answer
- 5.9% for Other with answer like :
 - We have management committee reviews and users committee reviews each 3 months

For a total of 17 complete answers and 1 skipped

2. Are there any additional processes should be defined in the PETAFL0P

environment?

Only one answer received (17 skipped):

- Public Relations, Dissemination

3. What are the main parts of your organisation?

management board

peer review team

administrators team

support & training team

dissemination team

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 76.5% for the Management board
- 47.1% for the Peer review team
- 88.2% for the administrator team
- 70.6% for the support and training team
- 23.5% for the dissemination team
- 35.3% for the Other with answers like:
 - Research teams
 - Networking
 - Developers, project workers
 - System and operations team

For a total of 17 complete answers and 1 skipped

*Management and decision support***1. How often do you perform procurement actions?**

Once a year

Monthly

At already predefined times

When needed

Other (please specify)

The answers received showed the following results

- 11% for the "Once a year" answer
- 0% for the Monthly answer
- 5% for the third answer
- 80% for the When needed answer
- 5% for the Other with answers like:
 - Typically we would procure a new system every two to three years, with upgrade options in the contracts to avoid the costly business of the procurement process.
 - HPC systems once a year

For a total of 17 complete answers and 1 skipped

2. Do you coordinate procurement of new systems among sites?

No

Yes, on university side level

Yes, on regional level

Yes, on national level

Who is responsible for performing the actions?

As answers were not exclusive, the answers received showed the following results

- 47.1% for the No answer
- 11.8% for the “Yes, on university side level” answer
- 0% for the “Yes, on regional level” answer
- 41.2% for the “Yes, on national level” answer
- 41.2% for the “Who is responsible for performing the actions?” with answers like:
 - I coordinate procurement only on CERFACS site
 - Sometimes during project, coordinator of the project is usually responsible
 - Handled by research councils
 - UNINETT Sigma and HPC/data centres
 - Gauss Centre for Supercomputing
 - Management committee for budget and planning, computing centre for technical specifications, analysis, and technical choice.
 - CSC
 - GENCI

For a total of 17 complete answers and 1 skipped

3. To whom do you need to report on?

Usage of your computing infrastructure

Scientific quality of computation results

As answers were not exclusive, the answers received showed the following results

- 100% for Usage of your computing infrastructure with answers like:
 - CERFACS' director
 - Advisory Board and management board of parent institution (ETHZ).
 - SNIC (national centre)
 - Ministry of Science and Higher Education
 - Ministry
 - Management board
 - Research councils
 - Research Council of Norway
 - Steering Committee and Commission for Informatics
 - CNRS direction
 - Polish Ministry of Science and Higher Education

For a total of 17 complete answers and 1 skipped

4. How often do you need to report on: usage of your computing infrastructure (i.e. system utilisation) or scientific quality of computation results, other?

Monthly

Yearly

Not at all

On demand

Other (please specify)

As answers were not exclusive, the answers received showed the following results

- 17.6% for the Monthly answer
- 70.6% for the Yearly answer
- 0% for the “Not at all” answer
- 41.2% for the “On demand” answer
- 20% for Other with answers like:
 - Quarterly usage statistics are provided in addition to annual report.
 - Every 6 months
 - 2x per year (I had to pick one of the answers here, to avoid error)
 - Each 3 months for usage of computing infrastructure and each year for scientific results

For a total of 17 complete answers and 1 skipped

5. Are you given the fixed report template document?

Yes

No

The answers received showed the following results

- 29.4% for the Yes answer
- 70.6% for the No answer

For a total of 17 complete answers and 1 skipped