



# notur

## User Survey

2009

---

**Title:** Notur User Survey 2009

**Version:** 1.0

**Release Status:** Final

**Availability:** Unrestricted

# FOREWORD

This report contains the results of the survey that was carried out in November 2009 - January 2010 among the research projects that make use of the national High Performance Computing facilities provided by the Notur project. The survey is repeated every year in more or less the same form. The aim of the survey is to determine overall user satisfaction with the provided resources and services by the Notur project. It is important for the project to provide an infrastructure that is in accordance with the needs and preferences of the users. This survey gives the user community a good opportunity to communicate their satisfaction with the resources and services, point to areas that need increased attention and inform the project about future needs for infrastructure. The results of this survey will be used to improve the operations and support mechanisms of the Notur project and guide the further planning of the project.

With respect to available resources in Notur, not much changed since the survey that was carried more than one year ago. The four computing resources that were used in the survey in 2008 are still in operation today in nearly the same configuration. The resources have now been available for general usage for 2-3 years and many users have built considerable experience in using these resources. Nevertheless, users develop or start using new applications, new users get access to the resources, and the resources undergo upgrades regularly. It is therefore important that a satisfactory level of user support is maintained for each system.

The key findings of the survey are presented on pages 3-6.

An electronic version of this report is available on [www.notur.no](http://www.notur.no). Reports with the results of the surveys from previous years can also be found there.

*Jacko Koster*  
*UNINETT Sigma AS*

## Contents

<b>1</b>	<b>About the Survey</b>	<b>1</b>
<b>2</b>	<b>Summary of the Survey</b>	<b>3</b>
<b>3</b>	<b>The Notur II project</b>	<b>6</b>
3.1	Notur Hardware Resources . . . . .	7
<b>4</b>	<b>User and Application Information</b>	<b>8</b>
4.1	User Information . . . . .	8
4.2	Application Information . . . . .	12
<b>5</b>	<b>Notur Hardware Resources</b>	<b>16</b>
<b>6</b>	<b>Notur Software Resources and Application Performance</b>	<b>21</b>
6.1	Notur Software Resources . . . . .	21
6.2	Application Performance . . . . .	22
<b>7</b>	<b>Notur User Support Services and Resource Allocation</b>	<b>27</b>
7.1	Notur User Support Services . . . . .	27
7.2	Advanced User Support . . . . .	28
7.3	Resource Allocation . . . . .	29
<b>8</b>	<b>Overall Satisfaction with Notur</b>	<b>31</b>
<b>9</b>	<b>Summary of Experiences</b>	<b>34</b>
<b>10</b>	<b>Future Needs</b>	<b>39</b>
<b>A</b>	<b>APPENDIX for Section 5</b>	<b>48</b>
<b>B</b>	<b>APPENDIX for Section 6.1</b>	<b>52</b>
<b>C</b>	<b>APPENDIX for Section 6.2</b>	<b>54</b>
<b>D</b>	<b>APPENDIX for Section 7.1</b>	<b>55</b>
<b>E</b>	<b>APPENDIX: Questionnaire 2009</b>	<b>57</b>

## List of Tables

1	Notur hardware resources (configuration November 2009) . . . . .	7
2	Home institution of the respondents . . . . .	9
3	”Which Notur resources did your project use in the last 12 months?” . . . . .	10
4	”How important is it for your research to have access to HPC resources?” . . . . .	11
5	”How long has your project used Notur resources?” . . . . .	11
6	”What is the frequency of your use of Notur resources in your project?” . . . . .	11
7	”Field of Science” . . . . .	11
8	Characterization of the applications . . . . .	12
9	Aspects important for the applications . . . . .	13
10	”Has your group recently performed measurements to determine how efficient your applications run on the Notur facilities, relatively to the peak performance of the machine/processors?” . . . . .	14
11	”Has your group recently optimized (fine-tuned) the performance of your applications on the Notur facilities?” . . . . .	14
12	”Are your most important applications suitable for ’capacity’ platforms? Capacity platforms provide high job throughput and can for example be clusters of compute nodes (each with 2-4 processors, 2-8 GB memory, 80-200 GB disk) that are connected with a fast interconnect (Myrinet, InfiniBand).” . . . . .	14
13	”Do you have any application that can only run (efficiently) on ’capability’ architectures? A capability machine is for example a shared-memory machine with a large number of processors that share a large pool of memory and where the processors are tightly connected through a high-performance (proprietary) interconnect.” . . . . .	14
14	”If you have answered yes to question 13, state the reasons why your application(s) require such platform.” . . . . .	14
15	”Answer only if you have answered yes to question 13. Is it possible to change the software with reasonable effort to make the application(s) run efficiently on capacity (cluster) machines? Please provide as much as possible detail.” . . . . .	15
16	Satisfaction with hardware resources (1) . . . . .	17
17	Satisfaction with hardware resources (2) . . . . .	18
18	”Please give additional comments on the various resources (possibly including those that you do not currently use). For instance, you may want to comment on how well the resources match (or do not match) your needs, and why. Please specify the resource in case of any machine-specific remarks. Specify improvements that you believe are necessary as well as details about any ’Very Dissatisfied’ answers given.” . . . . .	19
19	Satisfaction with software resources (1) . . . . .	22
20	Satisfaction with software resources (2) . . . . .	22

21	Satisfaction with application performance (1) . . . . .	23
22	Satisfaction with application performance (2) . . . . .	23
23	Maximum number of processors used per job ; maximum number of processors the application(s) can effectively use per job. . . . .	24
24	Ideal resource for the projects in 2010-2011. . . . .	25
25	”Specify improvements that you believe are necessary as well as details about any ‘Very Dissatisfied’ answers given. Please specify the resource or site in case of any machine-specific remarks. Also specify software that you believe Notur should offer but currently is missing.” . . . . .	26
26	Satisfaction with user support services (1) . . . . .	28
27	Satisfaction with support services (2) . . . . .	28
28	”Are you familiar with the advanced user support services provided by Notur?” .	29
29	”In case you have used it, were you satisfied with the results?” . . . . .	29
30	”Please specify other services that advanced user support should include that will suit your needs in the best possible way.” . . . . .	29
31	Satisfaction with Resource Allocation . . . . .	30
32	”Specify improvements that you believe are necessary as well as details about any ‘Very Dissatisfied’ answers given. Please specify the site in case of any machine- or site-specific remarks.” . . . . .	30
33	Overall satisfaction . . . . .	32
34	Overall importance . . . . .	33
35	”Specify improvements that you believe are necessary as well as details about any ‘Very Dissatisfied’ answers given. Please specify the site in case of any machine- or site-specific remarks.” . . . . .	33
36	Question 1. ”What does Notur do well? What aspects of Notur are you most pleased with?” . . . . .	34
37	Question 2. ”What should Notur do differently? How can Notur improve? Are there specific issues where Notur can do better (services, support, etc.)?” . . . . .	36
38	Question 3. ”How does Notur compare to other centers that you have experience with? What do other centers provide that Notur should provide as well? Please list the center(s) that you use in the comparison.” . . . . .	37
39	Question 1. ”Please give a brief description of any known changes in your research activity in the coming 24 months that will affect your needs for computational and storage resources and user support. This could for example be the termination of projects, start of new projects, specific deadlines, etc.” . . . . .	39
40	Question 2. ”Will the need for computational resources change significantly in your project in the next 24 months? If yes, please state the reason(s) and detail the changes in your needs.” . . . . .	42

41	Question 3. "Will the need for storage change significantly in your project in the next 24 months? If yes, please state the reason(s) and detail the changes in your needs." . . . . .	43
42	Question 4. "Will your project need access to hardware resources in 2009-2011 that are currently not provided by Notur (e.g., other computing architecture, other type of storage, etc.)? Please detail your needs." . . . . .	45
43	Question 5. "Will your project have a need for services in 2009-2011 that are currently not provided by Notur? Here, one could think of for example other types of support, other types of resource allocation, services for visualization, grid-enabling of applications, etc. Please detail your needs." . . . . .	46
44	Hardware_Uptime . . . . .	48
45	Hardware_JobWait . . . . .	48
46	Hardware_JobQueues . . . . .	48
47	Hardware_JobFlexibility . . . . .	49
48	Hardware_JobInteractive . . . . .	49
49	Hardware_DiskPermanent . . . . .	49
50	Hardware_DiskTemporary . . . . .	49
51	Hardware_DiskPerformance . . . . .	50
52	Hardware_TapeStorage . . . . .	50
53	Hardware_TapeResponse . . . . .	50
54	Hardware_TapeInterfaces . . . . .	50
55	Hardware_NetworkSpeed . . . . .	51
56	Hardware_Satisfaction . . . . .	51
57	Software_Environment . . . . .	52
58	Software_Compilers . . . . .	52
59	Software_Applications . . . . .	52
60	Software_Libraries . . . . .	53
61	Software_Perftools . . . . .	53
62	Software_Tools . . . . .	53
63	Application_Performance . . . . .	54
64	Application_SinglePerformance . . . . .	54
65	Application_ParallelPerformance . . . . .	54
66	Support_TechAdvice . . . . .	55
67	Support_ResponseTime . . . . .	55
68	Support_ResponseUrgent . . . . .	55
69	Support_Followup . . . . .	56

70	Support_TimeToSolve . . . . .	56
71	Support_Accounts . . . . .	56
72	Support_Pages . . . . .	56

# 1 About the Survey

The annual user survey for the Notur project was carried out in January 2010. The survey was similar to the surveys from 2006-2008. The reports from these previous surveys can be found on the Notur web pages.<sup>1</sup> In this section, we summarize the key aspects of the questionnaire. Appendix E includes the full questionnaire that was used.

A web-based questionnaire was used that contained both open and closed questions. The questionnaire was divided in seven sections:

- User and Application information
- Notur Hardware resources
- Notur Software resources and Application performance
- Notur User support services and Resource allocation
- Overall satisfaction with Notur
- Summary of experiences
- Future needs

The aim was to receive one completed questionnaire for each of the projects that were granted resources (CPU-hours) in allocation periods

- 2008.2 (October 2008 - March 2009), and
- 2009.1 (April 2009 - September 2009).

The project responsables were asked (by email) to complete the questionnaire, but could assign this task to a user of the project who has a good overview of the experiences and opinions of the other users in the project. Participation to the survey was voluntary.

The purpose of the survey and intended use of the results were stated in the introduction to the questionnaire. Different types of questions were included in the questionnaire. For most of the questions, the projects could answer by picking (at most) one of the provided options. The most frequently used type of questions are the following:

- Questions about user satisfaction with aspects of the Notur resources or services could be answered with one of the following values:
  - 'Very Satisfied' (VERYSAT) [5]
  - 'Satisfied' (SAT) [4]
  - 'Neutral' (NEUT) [3]
  - 'Dissatisfied' (DIS) [2]
  - 'Very Dissatisfied' (VERYDIS) [1]
  - 'No Answer' (NA) - default answer

The uppercase abbreviations are used in the remainder of this report. The values 1 to 5 enable us to perform some overall grading (averaging) of the answers to each question (but excluding the NAs).

---

<sup>1</sup><http://www.notur.no/publications>

- Questions regarding importance of particular Notur aspects could be answered with one of the following values:
  - 'Very Important' (VERYIMP) [3]
  - '(Somewhat) Important' (IMP) [2]
  - 'Not Important' (NOTIMP) [1]
  - 'No Answer' (NA) - default answer
- A number of questions could be answered with:
  - 'Yes'
  - 'No'
  - 'Do Not Know' - default answer
- The projects were also asked to provide comments in a text field. In particular, the projects were asked to provide additional details if they had answered 'Very Dissatisfied' to any of the questions.

In the remainder of the report, keywords (or abbreviations) are used in the table that are defined at the beginning of each section. This was primarily done to reduce the complexity and widths of the tables.

Table columns with header 'AVE' report the average ratings computed from the provided answers about satisfaction or importance (but excluding the NA answers). For example, if  $n_5$  projects answered 'Very Satisfied',  $n_4$  projects answered 'Satisfied', etc, then the average is computed as

$$(n_5 \cdot 5 + n_4 \cdot 4 + n_3 \cdot 3 + n_2 \cdot 2 + n_1 \cdot 1) / (n_5 + n_4 + n_3 + n_2 + n_1)$$

The lowest possible mark for satisfaction is 1, while the highest possible mark is 5. Similarly, the lowest possible mark related to importance is 1, while the highest possible mark is 3.

All answers were treated equally. That is, no distinction is made between answers from different scientific disciplines or organizations. In addition, no distinction is made between answers from projects with different needs (allocations) or different numbers of users.

We have to a large extent separated the presentation of the results from the observations, analysis and conclusions. Most of the observations and conclusions are given in Section 2.

All answers were treated confidentially. This was stated explicitly in the introduction text of the questionnaire. No identities are explicitly revealed in the publication of the results in this report. We did remove some explicit references to (CPU-hour) project account numbers. We did not remove explicit references to (self-developed) application software.

For open questions, the report lists the answers per project. The projects have been numbered for this purpose.

We have included the comments from the respondents where appropriate. The comments are not changed. However, spelling and grammar mistakes are corrected.

To help determining differences with the results from the survey from 2008, this report sometimes uses the format XX [YY], where XX is the result for this (2009) survey and YY the corresponding number from the 2008 survey.

## 2 Summary of the Survey

This section contains the main observations and conclusions of the survey. A brief overview of the Notur project is given in Section 3. This includes a summary of the hardware resources and mechanisms for resource allocation. The key aspects of the questionnaire and survey are summarized in Section 1. Detailed statistics and feedback from the users that responded to the questionnaire can be found in Sections 4-10 and Appendices A-D. Appendix E contains the full questionnaire that was used this year.

To illustrate differences with the results from 2008, we sometimes use the format XX [YY], where XX is the result for this (2009) survey and YY the corresponding number from the previous (2008) survey.

The results are extracted from 75 [52] questionnaires that were returned by projects that had access to the Notur supercomputer facilities in the period October 2008 - September 2009. The answers represent 77% [60%] of the 97 [87] projects that were contacted.

### User and Application Information

Half (51% [48%]) of the projects that replied have used Notur resources for more than 3 years and 80% [54%] of the projects use the Notur resources on a daily or weekly basis. As many as 88% [96%] of the respondents report that access to HPC resources is absolutely vital or very important for their research. The questionnaires were returned by projects connected to 17 [14] different organizations. More than 61% [61%] of the answers came from NTNU, UiB and UiO. The projects constitute 7 [7] different scientific disciplines, and 62% [75%] of the projects are related to chemistry, physics, material science or geosciences.

Nearly all the respondents (95% [96%]) consider their applications time demanding, while also many of the respondents characterize their applications as memory demanding (63% [71%]) and storage demanding (49% [63%]).

Similar to last year, high machine availability and reliability and high overall processing speed are considered most important. Also important are memory size, fast processor interconnect, large temporary storage capacity (at run-time) and large storage capacity for input and output data, in-depth user-support and the availability of programming tools and libraries.

About 41% [42%] of the respondents state that their most important applications are suitable for capacity platforms, while 16% [15%] state that they have applications that can only run efficiently on capability platforms.

About 32% [40%] of the projects has not recently performed measurements to determine the run-time efficiency of their applications or has tuned the performance of the applications. (Not all projects would be able to do so, since projects may not have access to the source code, e.g., in case of proprietary/commercial software, or may not have the necessary skills or time to do such measurements or tuning.)

### Hardware Resources

The respondents were in 2009 in general satisfied with the various aspects of the compute facilities. The projects report high satisfaction with the availability (uptime) of the resources, except for titan. (See Table 16.)

The high load on njord in 2009 has been a matter of concern to the respondents. This is reflected in relatively low satisfaction concerning waiting times for batch jobs. The respondents also note

the (limited) ability to run jobs interactively on njord. (See Table 16.)

Similar to last year, relatively few projects answered the questions regarding archiving / tape storage. This may be due to the fact that tape storage may not be of immediate importance to many of the projects or the respondents had insufficient knowledge or information about this aspect.

### **Software Resources and Application Performance**

The projects seem mostly satisfied with the software environments and compilers on the Notur facilities. The respondents report relatively low satisfaction concerning the performance analysis and debugging tools on titan. (See Table 19.)

The large majority of respondents (90%) report to be satisfied or very satisfied with the overall application performance on the Notur facilities. (This was 91% in 2008 and 62% in 2007.) In particular application performance on hexagon is very satisfactory to the respondents. (See Table 21.)

### **User Support and Resource Allocation**

The projects report in general high satisfaction with technical advice, response times (including those for urgent requests), follow-up to queries and reported problems, overall time to resolve problems and management of project and user accounts. Similar to 2008, the scores for titan are lower than for the other systems, in particular concerning response time and time to resolution. (See Table 26.)

The quality of the support pages and other end-user documentation has been a matter of concern issue in previous surveys. This year's results show that the perception of the users has not changed substantially and that there is a need (for *all* systems) to increase the effort to provide better information to the end-user and improve the present information. (See Table 26.)

The comments from the users concerning advanced user support indicate that there is an increasing need for longer-term (multi-year) support projects. (See Table 30.)

The feedback concerning the final allocations for period 2009.2 (starting October 1) was provided late to a number of large projects. (See Table 32.)

### **Overall Satisfaction**

The projects are satisfied to very satisfied with the Notur project as a whole. A number of projects are dissatisfied with the provided information and documentation, including web pages. There is high satisfaction with the provided hardware, the overall support organization and support services. (See Table 33.)

### **Summary of Experiences**

The areas where the projects are most satisfied are:

- User (help-desk) support
- Availability, reliability, stability, uptime
- Provided resources (in particular, total capacity)

Areas where improvements are necessary:

- Support web pages and available end-user documentation
- Performance analysis tools and machine availability (titan)

The need for archiving services (or persistent storage) should be covered by the NorStore project. According to the respondents, the services provided by the Notur project compare well to the services provided by HPC centres in other countries.

### **Future Needs**

Not surprisingly, a considerable number of projects expect an increase in the need for resources in the next 24 months, especially computational resources. The projects were not asked to quantify this in detail. Increased need for computational resources is generally the result of the need to run more and/or larger models, increase in staffing, or entering into a new project phase where compute-intensive software is taken into use. The respondents do not report an urgent need for new or other services that should be provided by the Notur project, though visualization is mentioned by some projects. (See Section 10.)

### **Comparison with Previous Surveys**

Similar surveys were carried out in 2006-2008. The questionnaire from 2008 was used as a basis for the questionnaire this year and only few minor changes were made.

The response rate for this survey (77%) is considerably higher than for the survey last year (60%). The absolute number of respondents is also higher (75 answers in 2009 versus 52 answers in 2008).

Overall, the responses in the 2009 survey are similar to the responses in the survey from 2008. Similarities include:

- Satisfaction with various aspects of hardware, software and (help-desk) support services
- General information, technical support pages and other end-user documentation must be improved

There is relatively low satisfaction concerning the uptime and stability of titan in 2009.

### **Action Items**

The Notur partners must study the results of the survey and address the feedback. The remarks must be checked to see whether they contain valuable information and require follow-up. The feedback shall be used to verify whether users need to be informed better or whether specific problems should be solved (or worked around).

The results from the survey show that the users are in general satisfied with the provided help-desk (email) support, but the provided technical support pages and end-user documentation needs further attention. This concerns the Notur project as a whole, and in particular the technical support pages for the facilities. Support pages with necessary and sufficient must exist for all facilities. Users are likely to encounter problems when porting their applications to environments with new or different software stacks. Providing information and documentation for new systems (or after major upgrades) is particularly important as it usually takes longer to change the perception/opinion of a user about the provided support than to change the contents of support pages. Some feedback in the survey from this year may be based on earlier experiences from the respondents, while web information may have improved in the past twelve months, but this merely underlines that information and documentation must be complete and up-to-date throughout the lifetime of a system.

It must be the ambition of the Notur project that the results of the next survey (that will be carried

out later in 2010) show a real increase in user satisfaction concerning the support pages and other documentation.

### Next User Survey

The Notur project will again carry out a survey late(r) in 2010. This survey will be similar to the survey that is reported in this document, but there will be added focus on user support and training.

The results from the next survey will be compared with the results from previous surveys. The list of project managers remains relatively constant over the years (many projects exist for many years) which allows us to monitor changes in user satisfaction and perception. When the next survey will be carried out, the computing resources will be similar to 2009.

## 3 The Notur II project

The Notur II project (or short Notur) provides the national infrastructure for High Performance Computing (HPC) and computational science in Norway.<sup>2</sup> A main goal of the project is to provide a modern, powerful, and cost-effective national infrastructure and enable efficient utilization of this infrastructure. The project serves the Norwegian computational science community by providing the infrastructure to individuals or groups involved in

- education and research at Norwegian universities, university colleges and research organizations
- operational forecasting at the Meteorological institute (until 12/2009)
- research and engineering at other organizations that collaborate with the Notur project

Notur is a project under the eVITA programme (eVitenskap, Infrastruktur, Teori og Anvendelser) from the Research Council of Norway (RCN).<sup>3</sup> UNINETT Sigma is the coordinating partner of the project. The consortium partners are (2009):

- NTNU - Norwegian University of Science and Technology, Trondheim
- UiB - University of Bergen
- UiO - University of Oslo
- UiT - University of Tromsø
- met.no - Meteorological Institute

The HPC centres at the four university partners host and operate the supercomputer facilities. Each facility typically consists of a computing part (multi-processor machine), a storage part (disk), and secondary storage for back-up (and sometimes archiving). The partners provide basic (helpdesk) support to the users by answering queries and helping to resolve problems with installing and running software, and management of data. The Notur project also provides higher-level (advanced) application support that aims at improving the utilization of the resources (e.g., by performance analysis, optimization, tuning and parallelization), by enabling complex applications and by providing new services to experienced as well as inexperienced users of the facilities. The partners also maintain a portfolio of (commercial or publicly available) software for a variety of scientific

---

<sup>2</sup>Notur: [www.notur.no](http://www.notur.no)

<sup>3</sup>eVITA: [www.rcn.no/evita](http://www.rcn.no/evita)

disciplines. A significant part of the users runs self-developed or community-developed software on the facilities.

Research groups/projects that are financed through the RCN or Ministry of Education and Research (KD) can apply for resources (CPU-hours) on one or more facilities for up to four 6-month (allocation) periods, starting April 1 or October 1. Applications for allocations are received by UNINETT Sigma and are evaluated by a Resource Allocation Committee (RFK) that is appointed by the RCN. The requested allocations in an applications may be modified by the RFK. Parameters in this process are for example the total amount of allocations requested versus the total available capacity, the contents of previous applications, scientific output and the usage of previous allocations by the applicant or research group. Once the application has been approved, the applicant will be responsible for a (CPU-hour) project during the allocation period. A project gets a certain number of CPU-hours (allocation) on one or more of the facilities. There are currently no restrictions set by the RFK for disk usage, but the HPC centres have local policies for this. The allocation for a project can be shared by several users. The project responsible decides who can use the allocation. A new user needs to submit an account application form to UNINETT Sigma that needs to be signed by the project responsible. (A user can belong to several projects.) Finally, if a project uses its entire allocation before the end of the period, the project responsible can apply for extra allocations.

### 3.1 Notur Hardware Resources

Table 1 lists some characteristics of the HPC facilities operated by the university partners of the Notur project.

Table 1: Notur hardware resources (configuration November 2009)

Name	System	Type	Number of usable cores	CPU type	Theoretical peak total (Tflop)	Total memory (GB)	Total disk capacity (TB)
hexagon	Cray XT4	MPP	5552	Opteron	51	6064	288
njord	IBM p575	SMP	944	power5+	7.5	2272	120
stallo	HP BL 460c	cluster	5632	Xeon	60	12064	128
titan	Sun x2200	cluster	2528	Opteron	ca. 23	1100	-

A brief description of the resources is given below. Hardware and software configurations of the resources change over time. The descriptions below cover the configurations from November 2009.

#### Hexagon

hexagon.bccs.uib.no is a Cray XT4 with 1388 nodes. Each node has a 2.3 Ghz quad-core AMD Opteron processor and 4 GB memory (128 nodes have 8 GB memory). The machine is intended for large scale parallel (MPI) applications with a tight coupling between the processes. The nodes are interconnected with Cray's SeaStar2 3D torus interconnect. Hexagon was taken into operation in January 2008 and made available for general use in February 2008. The last major upgrade took place in June 2008. Operating system: Linux. Site: UiB.

#### Njord

njord.hpc.ntnu.no is a cluster with 65 p575 nodes, each with 16 processors. The nodes have 32 GB

memory (four nodes have 128 GB memory). The machine is suitable for OpenMP applications that require at most 16 processors and large scale MPI applications with a tight coupling between the processes. The nodes are interconnected with IBM's HPS interconnect. Due to the memory used by the operating system in each node and the static memory partitioning implemented for the time-critical weather forecasting models from met.no, only 13 of the 32 GB memory per node is available for researchers on many of the nodes. Njord was taken into operation in November 2006 (and tripled in size December 2009). Operating system: AIX. Site: NTNU.

### **Stallo**

stallo.uit.no is an HP cluster with 704 nodes. Each node has two quad-core 2.66 Ghz Xeon processors, 16 GB memory and 120 GB disk (50 nodes have 32 GB memory). The system also has 128 TB centralized storage. The system is well-suited for most MPI applications and can be used for shared-memory OpenMP applications using at most eight processor cores. Ca. 60% of the nodes are connected with InfiniBand interconnect, the rest with standard Gigabit Ethernet. Stallo is available for general usage since January 2008. Operating system: Rocks Linux. Site: UiT.

### **Titan**

titan.uio.no is a cluster with Sun x2200 and x4600 nodes. The cluster was included in the Notur system October 2007. The Notur part of titan consists of 224 Sun x2200 nodes, each with two quad-core AMD Opteron processors, 16 GB memory and 500 GB disk. Also 160 existing nodes were upgraded to the same configuration. Three Sun x4600 with eight dual-core AMD Opteron processors (two with 128 GB memory, one with 256 GB memory) were added in March 2008. The system has InfiniBand interconnect between the nodes. The machine is well-suited for most MPI applications. The machine can be used for shared-memory OpenMP applications using at most eight cores. The Notur part of titan has a total of 2528 processor cores. Operating system: Rocks Linux. Site: UiO.

## **4 User and Application Information**

The first section of the questionnaire consisted of two parts. First, some details about the user/project were asked, followed by a set of questions that characterize the applications used by the project and what services the project considers important (for the applications). The section also included some questions related to performance measurements and performance tuning on the Notur resources.

### **4.1 User Information**

The participant was first asked to identify him-/herself by providing name and account number of the (CPU-hour) project. This information is not listed in this report. The home institution of the project (responsible) was derived from the project account information. A summary is given in Table 2.

The projects were asked to list the Notur resources they have used in the last 12 months. They could choose from the four hardware resources (Table 1), the Notur support services and the Notur web pages. Table 3 summarizes the answers per project.

Table 3 shows that some systems are mentioned more often than other systems. This should be taken into account when interpreting the results in later sections. E.g., the number of remarks

Table 2: Home institution of the respondents

institution	#
NTNU	20
UiB	7
UiO	19
UiT	3
UMB	2
HiO	1
HiMOLDE	1
met.no	8
NERSC	1
NIVA	1
SINTEF	4
Simula	1
IFE	3
NORSAR	1
NHH	1
Rikshospital	1
NumericalRocks	1
TOTAL	75

about a resource may be proportional to the number of respondents that use that resource, etc.

Table 3 also shows that the respondents refer both to (email) support and to the support web pages.

Table 3: "Which Notur resources did your project use in the last 12 months?"

nr	field	hexagon	njord	stallo	titan	support	web	total
1	Chemistry	-	-	x	x	-	x	3
2	Geosciences	x	x	x	-	x	x	5
3	Chemistry	x	x	x	-	-	x	4
4	Chemistry	-	-	-	x	-	x	2
5	Material Science	x	-	x	x	x	x	5
6	Physics	-	x	-	-	x	x	3
7	Physics	x	x	x	-	-	x	4
8	Geosciences	-	-	-	x	x	-	2
9	Physics	-	-	x	-	-	x	2
10	Physics	x	-	x	x	x	x	5
11	Mathematics and Informatics	x	-	x	x	x	x	5
12	Geosciences	-	x	-	x	x	-	3
13	Geosciences	-	-	x	-	-	-	1
14	Chemistry	-	x	x	-	-	-	2
15	Material Science	-	x	-	-	-	-	1
16	Computational Fluid Dynamics	-	x	x	-	x	-	3
17	Physics	-	-	-	x	x	x	3
18	Chemistry	-	-	-	x	x	-	2
19	Chemistry	-	-	x	-	x	x	3
20	Biosciences	-	-	x	-	x	-	2
21	Physics	x	-	-	x	x	x	4
22	Medical Science	-	-	x	-	x	x	3
23	Other	x	x	-	-	x	x	4
24	Geosciences	x	x	x	-	x	x	5
25	Physics	-	x	x	-	-	x	3
26	Physics	-	x	-	-	-	x	2
27	Other	-	-	x	x	x	x	4
28	Chemistry	-	x	x	-	x	x	4
29	Computational Fluid Dynamics	-	x	-	-	x	-	2
30	Geosciences	-	x	-	-	-	-	1
31	Geosciences	-	x	-	-	x	x	3
32	Geosciences	-	x	-	-	x	x	3
33	Geosciences	-	-	x	-	x	x	3
34	Computational Fluid Dynamics	x	x	-	-	x	x	4
35	Biosciences	x	-	-	-	-	-	1
36	Physics	-	-	x	x	-	-	2
37	Geosciences	-	-	x	x	x	x	4
38	Geosciences	x	-	-	x	x	-	3
39	Geosciences	-	x	-	x	-	x	3
40	Material Science	x	-	x	x	x	x	5
41	Other	-	-	-	x	x	-	2
42	Geosciences	-	-	-	x	-	x	2
43	Material Science	x	x	x	x	x	-	5
44	Mathematics and Informatics	-	-	x	-	x	x	3
45	Medical Science	-	-	x	-	x	x	3
46	Biosciences	-	-	x	x	x	x	4
47	Chemistry	x	x	x	x	x	x	6
48	Biosciences	-	-	x	x	-	-	2
49	Economics	-	-	-	x	x	x	3
50	Other	-	-	x	-	x	x	3
51	Computational Fluid Dynamics	-	x	-	-	-	-	1
52	Other	-	-	x	-	-	-	1
53	Chemistry	-	-	x	x	x	x	4
54	Chemistry	x	-	-	-	-	x	2
55	Geosciences	-	x	-	x	x	x	4
56	Geosciences	x	-	-	-	-	-	1
57	Geosciences	-	x	-	-	-	x	2
58	Physics	-	-	-	x	x	x	3
59	Biosciences	-	-	-	x	x	-	2
60	Physics	-	-	x	-	x	x	3
61	Petroleum Technology	-	-	x	-	-	-	1
62	Marine Technology	-	x	-	-	x	x	3
63	Material Science	-	x	-	-	x	-	2
64	Other	-	-	-	x	x	x	3
65	Biosciences	-	-	-	x	-	-	1
66	Material Science	x	-	x	-	-	x	3
67	Mathematics and Informatics	-	x	-	-	x	-	2
68	Biosciences	x	-	-	-	-	-	1
69	Mathematics and Informatics	-	-	-	x	x	-	2
70	Chemistry	x	-	-	-	-	x	2
71	Geosciences	x	x	-	x	x	x	5
72	Biosciences	-	-	-	x	x	x	3
73	Geosciences	-	-	-	x	-	x	2
74	Other	-	-	x	-	-	x	2
75	Marine Technology	-	x	x	-	x	x	4
total		21	29	36	33	47	49	
pct		28%	39%	48%	44%	63%	65%	

Table 4: "How important is it for your research to have access to HPC resources?"

ANSWER	#	pct
absolutely vital	49	65%
very important	17	23%
important	9	12%
useful	0	0%
not important at all	0	0%
no answer	0	0%
TOTAL	75	100%

Table 5: "How long has your project used Notur resources?"

ANSWER	#	pct
six months or less	5	7%
six months to three years	31	41%
more than three years	38	51%
do not know	1	1%
TOTAL	75	100%

Table 6: "What is the frequency of your use of Notur resources in your project?"

ANSWER	#	pct
daily	35	47%
weekly	25	33%
monthly	9	12%
less than monthly	3	4%
do not know	3	4%
TOTAL	75	100%

Table 7: "Field of Science"

ANSWER	#	pct
Biosciences	8	11%
Chemistry	11	15%
Computational Fluid Dynamics	4	5%
Geosciences	18	24%
Material Science	6	8%
Mathematics and Informatics	4	5%
Physics	11	15%
Structural Engineering	0	0%
Other	7	9%
No Answer	0	0%
TOTAL	75	100%

## 4.2 Application Information

The projects were asked to characterize their applications and state what they consider important. The following items could be chosen to characterize the applications:

Application_Time:	''Time (CPU-hour) demanding''
Application_Memory:	''Memory demanding''
Application_Storage:	''Storage demanding''
Application_Proprietary:	''Proprietary / commercial software''
Application_Public:	''Public domain software''
Application_Selfdeveloped:	''Self-developed software''
Application_Sequential:	''Sequential application''
Application_EmbarrassinglyParallel:	''Embarrassingly parallel application''
Application_SharedMemory:	''Shared-memory application (e.g., using OpenMP)''
Application_DistributedMemory:	''Distributed-memory application (e.g., using MPI)''

Table 8 summarizes the answers returned by the projects.

Table 8: Characterization of the applications

KEY	YES	NO	DON'T KNOW	TOTAL
Application_Time pct	71 95%	3 4%	1 1%	75 100%
Application_Memory pct	47 63%	21 28%	7 9%	75 100%
Application_Storage pct	37 49%	24 32%	14 19%	75 100%
Application_Proprietary pct	22 29%	43 57%	10 13%	75 100%
Application_Public pct	44 59%	22 29%	9 12%	75 100%
Application_Selfdeveloped pct	55 73%	15 20%	5 7%	75 100%
Application_Sequential pct	22 29%	28 37%	25 33%	75 100%
Application_EmbarrassinglyParallel pct	33 44%	22 29%	20 27%	75 100%
Application_SharedMemory pct	26 35%	24 32%	25 33%	75 100%
Application_DistributedMemory pct	36 48%	12 16%	27 36%	75 100%

Nearly all respondents 95% [96%] consider their applications time (CPU-hour) demanding, while 63% [71%] of the respondents characterize their applications as memory demanding. A project can use a combination of sequential, embarrassingly parallel, shared-memory parallel, and distributed-memory parallel applications. Similarly, projects can use a combination of proprietary/commercial, public domain, and self-developed software.

The following items could be chosen to state what the projects consider important:

Application_MachineAvailability:	"High machine availability and reliability"
Application_ProcessingSpeed:	"High overall processing speed"
Application_MemorySize:	"Large overall memory size"
Application_Interconnect:	"Processor interconnect: high bandwidth & low latency"
Application_Security:	"High system (data) security"
Application_TempStorage:	"Large temporary storage capacity at run-time"
Application_Storage:	"Large disk storage capacity for input and output data"
Application_TapeStorage:	"Large tape storage capacity for archiving"
Application_NetworkCapacity:	"High external network capacity to access the system"
Application_UserSupport:	"Availability of in-depth user support from Notur"
Application_ProgrammingTools:	"Availability of programming tools and libraries"
Application_PerformanceTools:	"Availability of performance tools and debuggers"

Table 9 summarizes the answers returned by the projects.

Table 9: Aspects important for the applications

KEY	YES	NO	DON'T KNOW	TOTAL
Application_MachineAvailability pct	67 89%	2 3%	6 8%	75 100%
Application_ProcessingSpeed pct	68 91%	2 3%	5 7%	75 100%
Application_MemorySize pct	43 57%	19 25%	13 17%	75 100%
Application_Interconnect pct	37 49%	12 16%	26 35%	75 100%
Application_Security pct	13 17%	39 52%	23 31%	75 100%
Application_TempStorage pct	39 52%	20 27%	16 21%	75 100%
Application_Storage pct	43 57%	20 27%	12 16%	75 100%
Application_TapeStorage pct	16 21%	38 51%	21 28%	75 100%
Application_NetworkCapacity pct	31 41%	22 29%	22 29%	75 100%
Application_UserSupport pct	41 55%	18 24%	16 21%	75 100%
Application_ProgrammingTools pct	39 52%	18 24%	18 24%	75 100%
Application_PerformanceTools pct	27 36%	24 32%	24 32%	75 100%

The table shows that many of the items are considered important by the projects, in particular high machine availability and reliability and high overall processing speed. Security and tape storage appear to be less important.

The projects were also asked questions about performance measurements and performance tuning of their applications and which type of hardware platforms they consider (most) suitable. The questions and answers are summarized in Tables 10-15.

Table 10: "Has your group recently performed measurements to determine how efficient your applications run on the Notur facilities, relatively to the peak performance of the machine/processors?"

KEY	YES	NO	DON'T KNOW	TOTAL
Application_Measure	24	45	6	75
pct	32%	60%	8%	100%

Table 11: "Has your group recently optimized (fine-tuned) the performance of your applications on the Notur facilities?"

KEY	YES	NO	DON'T KNOW	TOTAL
Application_Tune	22	50	3	75
pct	29%	67%	4%	100%

Table 12: "Are your most important applications suitable for 'capacity' platforms? Capacity platforms provide high job throughput and can for example be clusters of compute nodes (each with 2-4 processors, 2-8 GB memory, 80-200 GB disk) that are connected with a fast interconnect (Myrinet, InfiniBand)."

KEY	YES	NO	DON'T KNOW	TOTAL
Application_Capacity	31	8	36	75
pct	41%	11%	48%	100%

Table 13: "Do you have any application that can only run (efficiently) on 'capability' architectures? A capability machine is for example a shared-memory machine with a large number of processors that share a large pool of memory and where the processors are tightly connected through a high-performance (proprietary) interconnect."

KEY	YES	NO	DON'T KNOW	TOTAL
Application_Capability	12	32	31	75
pct	16%	43%	41%	100%

Table 14: "If you have answered yes to question 13, state the reasons why your application(s) require such platform."

Number of comments: 13

- 1: For the bulk of our applications, the answer to question 13 is "no". However, it is necessary for us to do high-accuracy calculations for calibration purposes, and these demand "capability" architectures.
- 2: One of our main applications, Bergen Climate Model (BCM), requires a shared memory machine to achieve a useable integration speed. The Norwegian Earth System Model (NorESM) is in its final stage of development

and will gradually replace BCM. NorESM can run efficiently on capacity platforms. However, since BCM experiments are required in several ongoing projects, efficient execution of BCM is still important for members of nnXXXXk.

- 3: The various electron-structure methods differ greatly with respect to scalability. Whereas density-functional theory methods run well on capacity platforms, most implementations of coupled-clusters methods (like CCSD(T) in gaussian) need to run in shared-memory mode.
- 5: Some of our jobs require very high memory per CPU, higher than what is commonly available on most of the NOTUR machines.
- 8: Paralellized with OpenMP.
- 19: Due to high I/O.
- 20: I am not sure what "capability" architectures means for our applications, therefore the answer is "Do not know".
- 22: Too time consuming on only few processors.
- 24: Our applications HYCOM and EnKF are applied to large 3D discrete grids where the grid size can typically be 1000x1000x30 with several 3D variables. The software has been modified to support MPI-parallelization and splitting of the grid into tiles, but the processors need to be tightly connected to pass messages between tiles.
- 31: We are running a fully coupled regional climate model (atmosphere-land, ice and ocean). The communication between the component models uses MPI. However, the ocean component has only a shared memory openMP parallelization internally. Therefore it has to run on 16-32 threads to keep in phase with the rest of the system.
- 34: One of our most used programs, SSIIM, is parallelized using OpenMP. It scales linearly on Njord, to 16 processors. It would be very useful for us to use a shared memory machine with more than 16 processors. It is not feasible to convert the program to use distributed memory machines/MPI.
- 47: There is an increasing activity in the group towards calculations using periodic-boundary condition plane-wave codes, largely through commercial codes that require fast interconnects and large memory in order to run efficiently.
- 75: It was developed at a time where these hw-architectures were available.

Table 15: "Answer only if you have answered yes to question 13. Is it possible to change the software with reasonable effort to make the application(s) run efficiently on capacity (cluster) machines? Please provide as much as possible detail."

Number of comments: 11

- 1: Not really. The complexity leading to the architecture requirements mostly reflect the underlying mathematical structure of the model.
- 2: Porting BCM to a capacity platform is demanding. One of the objectives for replacing BCM with NorESM is the ability to utilise capacity platforms. Therefore we believe it is not worthwhile to put a lot of effort into making BCM run efficiently on capacity platforms. Only a limited number of BCM experiments are promised in ongoing and future projects. Our hope is that we can be granted access to capability platforms for the remainder of BCM experiments.
- 3: With respect to gaussian: No. However, we are working toward replacements, and the situation may change during the coming year.
- 5: This is doubtful. When using the hybrid functional option of VASP, the calculations require very high memory due to the large matrices being diagonalized.
- 8: No. It takes too long on 2-4 processors.
- 19: With time it should be possible ... but will require many man hours.
- 24: (following above question) The messages passed between tiles correspond to ocean processes (advection of tracers, large-scale waves...) and need to be done at high frequency. On a cluster with limited interconnect, the scaling will be poor and there is little to do with the software to improve that.
- 31: In a few (1-2?) years, the openMP based ocean component will be replaced with a model parallelized with MPI. Then it should be able to use a capacity machine. Before this, the investment of rewriting the code for MPI will take too long to be practical within the project.

- 34: No. It took one year to parallelize the program using OpenMP. It will take at least ten times longer to do the same with MPI.
- 47: Since the currently used codes are commercial codes, it is hard for us to move these applications to a capacity machine. However, we are in the process of developing our in-house PBC code. However, it is difficult to state whether this code will be able to fully replace the commercial codes and/or the human effort required to provide a functionality comparable to that of the commercial packages and at the same time would run efficiently on capacity platforms. However, it is reasonable to assume this effort to amount to multiple PhD/postdoc positions.
- 75: We think so, but we need help - "reasonable effort" ??

## 5 Notur Hardware Resources

The projects were asked to rate their satisfaction with the Notur hardware resources that they have used. The following keywords are used in this section:

Hardware_Uptime:	"Machine availability (uptime)"
Hardware_JobWait:	"Batch job wait time"
Hardware_JobQueues:	"Batch queue structure and job priorities"
Hardware_JobFlexibility:	"Batch system flexibility"
Hardware_JobInteractive:	"Ability to run interactively"
Hardware_DiskPermanent:	"Available storage capacity for permanent data (home directories)"
Hardware_DiskTemporary:	"Available storage capacity for temporary data"
Hardware_DiskPerformance:	"File system configuration and I/O performance"
Hardware_TapeStorage:	"Available storage capacity for archiving"
Hardware_TapeResponse:	"Performance of archiving/retrieval tools"
Hardware_TapeInterfaces:	"User interfaces for archiving/retrieval of data"
Hardware_NetworkSpeed:	"External network capacity to access the system and transfer data"
Hardware_Satisfaction:	"Overall satisfaction"

The answers are summarized in Tables 16 and 17. The numbers in Table 16 represent the average rating and the percentage of the projects that have evaluated the aspect (keyword) for a particular resource. The numbers in Table 17 represent the number of projects that have given the same grade for a particular aspect (keyword) totalled over all resources. Hence, the numbers in column TOTAL can exceed the total number of responses. The average ratings in the columns AVE of Tables 16 and 17 are the same.

The average ratings in the columns AVE of Tables 16 and 17 show that the projects are in general satisfied with the various aspects of the hardware. The projects reported high satisfaction with the hardware availability (uptime) which they consider also important for their applications (see Table 9). The responses to tape issues are relatively few. In Appendix A, more detailed results are presented for each of the keywords and per resource. More observations can be found in Section 2.

Table 16: Satisfaction with hardware resources (1)

KEY	hexagon	njord	stallo	titan	AVE
Hardware_Uptime	4.5	4.4	4.5	3.5	4.2
pct	25%	35%	47%	41%	-
Hardware_JobWait	4.1	3.0	4.0	4.0	3.8
pct	25%	36%	44%	37%	-
Hardware_JobQueues	4.0	3.5	3.8	3.7	3.7
pct	23%	33%	41%	37%	-
Hardware_JobFlexibility	4.1	3.7	3.8	3.6	3.8
pct	21%	27%	35%	35%	-
Hardware_JobInteractive	3.9	3.2	3.8	3.6	3.6
pct	15%	24%	29%	25%	-
Hardware_DiskPermanent	3.8	3.6	4.0	3.7	3.8
pct	23%	33%	39%	35%	-
Hardware_DiskTemporary	4.5	4.0	4.0	3.9	4.1
pct	23%	33%	39%	29%	-
Hardware_DiskPerformance	4.3	4.1	3.7	3.8	3.9
pct	23%	29%	35%	33%	-
Hardware_TapeStorage	4.0	3.7	3.8	3.9	3.8
pct	19%	24%	24%	20%	-
Hardware_TapeResponse	3.5	3.7	3.6	3.4	3.6
pct	16%	19%	19%	12%	-
Hardware_TapeInterfaces	3.5	3.4	3.5	3.4	3.5
pct	16%	19%	19%	11%	-
Hardware_NetworkSpeed	4.1	4.1	4.1	4.0	4.1
pct	24%	28%	39%	32%	-
Hardware_Satisfaction	4.4	4.0	4.3	3.8	4.1
pct	24%	35%	47%	43%	-

Table 17: Satisfaction with hardware resources (2)

KEY	VERYSAT	SAT	NEUT	DIS	VERYDIS	TOTAL	AVE
Hardware_Uptime	51	42	13	1	4	111	4.2
pct	46%	38%	12%	1%	4%	100%	-
Hardware_JobWait	25	46	26	7	3	107	3.8
pct	23%	43%	24%	7%	3%	100%	-
Hardware_JobQueues	16	50	27	7	1	101	3.7
pct	16%	50%	27%	7%	1%	100%	-
Hardware_JobFlexibility	13	44	29	2	0	88	3.8
pct	15%	50%	33%	2%	0%	100%	-
Hardware_JobInteractive	7	39	16	7	1	70	3.6
pct	10%	56%	23%	10%	1%	100%	-
Hardware_DiskPermanent	26	34	30	5	2	97	3.8
pct	27%	35%	31%	5%	2%	100%	-
Hardware_DiskTemporary	28	47	15	2	1	93	4.1
pct	30%	51%	16%	2%	1%	100%	-
Hardware_DiskPerformance	24	39	24	3	0	90	3.9
pct	27%	43%	27%	3%	0%	100%	-
Hardware_TapeStorage	18	21	23	2	1	65	3.8
pct	28%	32%	35%	3%	2%	100%	-
Hardware_TapeResponse	7	15	26	1	0	49	3.6
pct	14%	31%	53%	2%	0%	100%	-
Hardware_TapeInterfaces	6	11	30	1	0	48	3.5
pct	12%	23%	62%	2%	0%	100%	-
Hardware_NetworkSpeed	27	45	20	0	0	92	4.1
pct	29%	49%	22%	0%	0%	100%	-
Hardware_Satisfaction	32	64	9	5	1	111	4.1
pct	29%	58%	8%	5%	1%	100%	-

Table 18: "Please give additional comments on the various resources (possibly including those that you do not currently use). For instance, you may want to comment on how well the resources match (or do not match) your needs, and why. Please specify the resource in case of any machine-specific remarks. Specify improvements that you believe are necessary as well as details about any 'Very Dissatisfied' answers given."

#### Resource hexagon. 7 comments:

- 10: Job priorities on hexagon work as long as the machine is not heavily overloaded but may not give fair priorities at times of heavy oversubscription.
- 21: We are generally pleased with hexagon. Things are running smoothly there.
- 35: Very happy with the performance on hexagon. For our jobs it is the only hpc facility in Norway to use.
- 40: Home directories not sufficient..
- 47: We have just started to use this machine for commercial application, and have thus limited experience with the use of this machine.
- 54: If the number of processors (cores) on one motherboard can be further increased, the efficiency of our application will be improved to a large extent.
- 68: We have made little use of hexagon so far. Unfortunately data collection was much more time consuming than we thought so we were not able make progress with analysis.

#### Resource njord. 9 comments:

- 2: We have used njord mainly for BCM experiments. These experiments only fit in the "normal" queue and some users have reported quite long waiting time. This is probably due to big load on the machine and is not a system problem as we see it. In addition to the experiments, sequential post processing applications must be executed. These can only fit in the "single" queue and users have reported congestion of this queue. We understand the machine is mainly for parallel applications but this post processing step is an essential process for us. Moving the large amount of output/restart data to another machine for post processing is not practical and we feel there should be enough flexibility in the queue system to enable efficient execution of these serial applications.
- 6: The xlc-compiler is excellent. We look forward to the next generation of machines.
- 12: Until the extension from 56 to 192 nodes, it was very difficult to run jobs that were not in the 'forecast' queue. We have not yet enough experience to tell if this situation has improved.
- 16: Better access for single processor jobs.
- 25: Could have been easier to get an overview of the usage on a given account.
- 30: Since we are producing a lot of output it would be nice if a better storage/archiving system would be available. I always have a problem to store the produced data.
- 34: We are very pleased with the installment of the new processors on Njord. The availability of the computer became very much better.
- 47: We have just started to use this machine for commercial application, and have thus limited experience with the use of this machine.
- 51: (1) I don't know whether there are some interactive nodes on njord, which is more convenient for code debugging and development. (2) The waiting time is long for batch jobs.

#### Resource stallo. 12 comments:

- 2: We have not used this platform enough to make a proper assessment.
- 10: We have had very little usage of stallo and have therefore not given detailed answers
- 13: The only real limitation in our applications is the number of available CPU. So far there has not been too much queue on Stallo and we hope this can continue. We often need to run a large number (up to several hundreds) of "scenarios" and get the results as fast as possible. In this sense we need much more CPU than what is needed for a single job.

- 25: Could have been easier to get an overview of the usage on a given account.
- 33: We are extremely satisfied with the resources available and the support of the notur team when we have asked for help. The relatively straightforward and painless application/reporting process for resources means we can spend our efforts on science rather than administration. The flexibility of allocation of resources means we have the opportunity to try new types of models and request additional resources if initial tests are successful rather than upfront specifying large resources that we may not in the end use. I guess the permanent storage question is not so applicable to us given that we have been given norstore resources.
- 40: Backup of temporary files would be appreciated
- 44: We are in general very satisfied with Stallo.
- 46: I think the resources are fairly well suited to our needs, though disk IO is a massive bottleneck for the storage strategies I've been using so far. Furthermore, we have been much delayed in our use of both embarrassingly parallel and massively parallel programs; we hope to do better next year. An option that might become highly useful in the future is interactive parallel work using IPython, <http://ipython.scipy.org/doc/rel-0.10/html/overview.html#interactive-parallel-computing>
- 48: We value stallo for its high availability and the fast reply times we have had from local user support in Tromsø. After solving some problems in our NEST software (sub-optimal use of MPI leading to deadlocks) we now see excellent scaling also for large numbers of processors. The questions on technical aspects asked above do not match our requirements very well, so I will sum up here: Our simulations are mainly limited by RAM available for storing network connectivity information. Our MPI-communication requires rather little bandwidth but is sensitive to latency.
- 50: I run ROMS on Stallo, and Stallo matches my needs.
- 53: I do not use stallo personally, but my impression is that things are working well.
- 60: With exception of some heavy test calculations, I have devoted most of my time to program development. For applications, the program requires several compute nodes with large memories, i.e greater than 16GB.

#### Resource titan. 14 comments:

- 8: Titan was very unstable in the first part of 2009, but seems to be more stable now.
- 10: We have had very little usage of titan and have therefore not given detailed answers
- 12: More easy to get into the run queue here. However, we have experienced some instability of this HPC system.
- 17: My biggest concern so far has been the various problems of stability we've seen from time to time. There have been several MPI problems, disk crashes etc. that have been problematic at times. This has generally become a lot better lately, and I do hope this trend will continue.
- 21: The problems with the upgrade in april of titan caused serious problems and it took a long time before things were fixed, in particular the new queue system had too many bugs. This caused a huge waste of cpu time. Files which were supposed to be stored after a run were never stored. Some of my students who had to finish their master theses in june had to migrate all their jobs to hexagon, simply because titan didn't work. If new upgrades are planned, one should avoid doing this in the months before students hand their theses. The standard time for handing in theses are now June 1 and December 1. The online FAQs lack important details and are not well organized.
- 40: Have not used this machine lately
- 46: I think the resources are fairly well suited to our needs, though disk IO is a massive bottleneck for the storage strategies I've been using so far. Furthermore, we have been much delayed in our use of both embarrassingly parallel and massively parallel programs; we hope to do better next year. An option that might become highly useful in the future is interactive parallel work using IPython, <http://ipython.scipy.org/doc/rel-0.10/html/overview.html#interactive-parallel-computing>
- 48: In comparison to stallo, titan appears less stable and more crowded. From our perspective, it would have been preferable to prioritize stability over expansion on titan. Ideally, we would like to see identical queuing systems used on stallo and titan: both machines are quite similar in their architecture so we in principle could move jobs easily between both machines. The different queing systems make this more demanding than we believe necessary.
- 49: I have tried to use the lowpri queue, but have not been able to figure out how to correctly submit the -ckpt blcr switch. The documentation is a bit sparse on this point.

- 53: A lot of crashes, probably related to the queuing system in 2009.
- 55: We have had and continues to have problems compiling Fortran based MPI codes on titan. These programs worked fine with MPICH2 libraries which were available in the earlier version of titan, but don't work with openMPI which are perhaps the only available ones in the present version of titan. It would be useful to have several flavours of MPI libraries to work with and test our programs rather than be constrained with just open source solutions like openMPI.
- 64: The three large-memory nodes tend to be over-subscribed, i.e. it can be difficult to put jobs on these nodes (with reasonable wait times). One of our applications has a memory footprint of around 32 gbytes, so for our purposes it would be desirable to have more medium-memory nodes (with 64 gbytes RAM, say) available.
- 72: I have not yet been able to actually make use of the Notur resources. We had applied for a shared project area and I only just heard that we can't actually get that (though they are working on it now!). I am therefore not yet able to answer the above questions. Overall the support we have been getting is good, though it was unclear that we didn't get the shared project area we applied for and has a result it has taken a very long time to get things setup. It would have been good if we had been told from the beginning more clearly what we could and could not get.
- 73: Overall we are very satisfied with titan in particular with respect to the amount of resources granted for this project, which are critical for the kind of work we do.

## 6 Notur Software Resources and Application Performance

This section of the questionnaire consisted of two parts. The projects were asked to state their satisfaction regarding the software resources provided by Notur and with the performance of their own (application) software.

### 6.1 Notur Software Resources

The projects were asked to rate their satisfaction with the software resources provided on each of the Notur hardware resources that they had used. The following keywords are used in this section:

Software_Environment:	"User environment (shells, etc.)"
Software_Compilers:	"Compilers"
Software_Applications:	"Application software"
Software_Libraries:	"Programming tools and libraries"
Software_Perftools:	"Performance and debugging tools"
Software_Tools:	"General tools and utilities"

The ratings in Tables 19 and 20 were calculated in the same way as those in Tables 16 and 17.

Table 19: Satisfaction with software resources (1)

KEY	hexagon	njord	stallo	titan	AVE
Software_Environment	4.4	4.1	4.2	4.1	4.2
pct	23%	35%	40%	40%	-
Software_Compilers	4.3	4.3	4.0	4.0	4.1
pct	20%	31%	33%	29%	-
Software_Applications	4.3	3.8	3.8	3.7	3.8
pct	12%	25%	28%	24%	-
Software_Libraries	4.2	3.8	3.8	3.9	3.9
pct	16%	24%	25%	23%	-
Software_Perftools	4.1	3.8	3.6	3.1	3.6
pct	12%	19%	24%	16%	-
Software_Tools	4.3	3.8	4.0	3.8	3.9
pct	16%	25%	29%	23%	-

Table 20: Satisfaction with software resources (2)

KEY	VERYSAT	SAT	NEUT	DIS	VERYDIS	TOTAL	AVE
Software_Environment	35	55	14	1	0	105	4.2
pct	33%	52%	13%	1%	0%	100%	-
Software_Compilers	27	47	11	2	0	87	4.1
pct	31%	54%	13%	2%	0%	100%	-
Software_Applications	11	38	18	2	0	69	3.8
pct	16%	55%	26%	3%	0%	100%	-
Software_Libraries	12	38	16	1	0	67	3.9
pct	18%	57%	24%	1%	0%	100%	-
Software_Perftools	8	22	18	6	0	54	3.6
pct	15%	41%	33%	11%	0%	100%	-
Software_Tools	11	44	16	0	0	71	3.9
pct	15%	62%	23%	0%	0%	100%	-

Satisfaction with the provided software resources is high. The category 'Performance and debugging tools' scores somewhat lower than the other.

In Appendix B, more detailed results are presented for each of the keywords per resource. More observations can be found in Section 2.

## 6.2 Application Performance

The projects were asked about the overall satisfaction with the performance of the applications used by the project. The answers and comments are summarized in Tables 21-25.

The following keywords are used in this section:

Application_Performance:	”Overall application performance”
Application_SinglePerformance:	”Sequential application performance”
Application_ParallelPerformance:	”Parallel application performance (scalability)”
Application_NrProcsUsed:	”Max. number of processors used per job in your project”
Application_NrProcsPossible:	”Max. number of processors your application can use per job (in ideal circumstances)”

The ratings in Tables 21 and 22 are computed in the same way as those in Tables 16 and 17.

Table 21: Satisfaction with application performance (1)

KEY	hexagon	njord	stallo	titan	AVE
Application_Performance	4.6	4.1	4.2	3.9	4.2
pct	19%	31%	39%	35%	-
Application_SinglePerformance	4.4	3.8	4.1	3.9	4.0
pct	7%	16%	20%	23%	-
Application_ParallelPerformance	4.6	4.1	4.2	3.8	4.1
pct	17%	28%	37%	28%	-

Table 22: Satisfaction with application performance (2)

KEY	VERYSAT	SAT	NEUT	DIS	VERYDIS	TOTAL	AVE
Application_Performance	26	58	8	2	0	94	4.1
pct	28%	62%	9%	2%	0%	100%	-
Application_SinglePerformance	12	27	11	1	0	51	4.0
pct	24%	53%	22%	2%	0%	100%	-
Application_ParallelPerformance	26	47	8	4	0	85	4.1
pct	31%	55%	9%	5%	0%	100%	-

As many as 90% [91%] of the respondents report to be very satisfied or satisfied with the overall application performance on the Notur facilities. (This was 62% in 2007.) In Appendix C, more detailed results are presented for each of the keywords per resource. More observations can be found in Section 2.

The projects were asked to state for each resource (i) the maximum number of processors used per job and (ii) the maximum number of processors their code(s) can use effectively per job. The reported numbers are listed in Table 23. For these questions, we list all the individual answers, as they are very diverse (instead of reporting averages).

The projects were also asked what would be the ideal resource for their research in 2010-2011. This consisted of giving five parameters for the resource that realistically can be used by the project. The results are shown in Table 24 for each project that answered the questionnaire. Also here, we list all the individual answers, instead of reporting averages. The projects that did not return answers are not listed. The following keywords are used:

IdealResource_CPUs:	”Number of CPUs”
IdealResource_MemoryConfig:	”Memory configuration”
IdealResource_Memory:	”Memory (GB per processor)”
IdealResource_Storage:	”Storage (TB)”
IdealResource_Interconnect:	”Processor Interconnect”

Table 23: Maximum number of processors used per job ; maximum number of processors the application(s) can effectively use per job.

project	hexagon	njord	stallo	titan
1	- ; -	- ; -	128 ; 1000	32 ; 128?
2	300 ; 1000	16 ; 32	0 ; 0	- ; -
3	- ; -	- ; -	- ; -	- ; -
4	- ; -	- ; -	- ; -	0 ; 0
5	32 ; 128	- ; -	32 ; 128	32 ; 128
6	- ; -	1 ; 1	- ; -	- ; -
7	200 ; -	32 ; -	200 ; -	- ; -
8	- ; -	- ; -	- ; -	8 ; 32
9	- ; -	- ; -	128 ; 256	- ; -
10	1024 ; >2000	- ; -	0 ; 0	0 ; 0
11	256 ; 1024	- ; -	1024 ; 4096	128 ; 1024
12	- ; -	128 ; 400	- ; -	64 ; 200
13	- ; -	- ; -	1024 ; 1024	- ; -
14	- ; -	16 ; 0	16 ; 0	- ; -
15	- ; -	64 ; 64	- ; -	- ; -
16	- ; -	32 ; 96	32 ; 96	- ; -
17	- ; -	- ; -	- ; -	400 ; 1000
18	- ; -	- ; -	- ; -	1 ; 0
19	- ; -	- ; -	1 ; 1	- ; -
20	- ; -	- ; -	8 ; 8	- ; -
21	1000 ; all if possible	- ; -	- ; -	1000 ; all if possible
22	- ; -	- ; -	64 ; ?	- ; -
23	200 ; 2000	20 ; 20	- ; -	- ; -
24	- ; -	- ; -	- ; -	- ; -
25	- ; -	32 ; 0	32 ; 0	- ; -
26	- ; -	unknown ; unknown	- ; -	- ; -
27	- ; -	- ; -	256 ; 0	32 ; 0
28	- ; -	256 ; 256	256 ; 256	- ; -
29	- ; -	16 ; 32	- ; -	- ; -
30	- ; -	16 ; 0	- ; -	- ; -
31	- ; -	64 ; 64	- ; -	- ; -
32	- ; -	32 ; 32	- ; -	- ; -
33	- ; -	- ; -	60 ; 256	- ; -
34	512 ; infinite	512 ; infinite	- ; -	- ; -
35	0 ; 0	- ; -	- ; -	- ; -
36	- ; -	- ; -	1 ; 1	1 ; 1
37	- ; -	- ; -	96 ; 600	64 ; 600
38	0 ; 0	- ; -	- ; -	0 ; 0
39	- ; -	64 ; 512	- ; -	64 ; 512
40	0 ; 0	- ; -	0 ; 0	0 ; 0
41	- ; -	- ; -	- ; -	1 ; 1
42	- ; -	- ; -	- ; -	64 ; 256
43	36 ; 36	16 ; 16	8 ; 8	32 ; 32
44	- ; -	- ; -	8 ; Depends on per-node memory	- ; -
45	- ; -	- ; -	320 ; unlimited	- ; -
46	- ; -	- ; -	8 ; 64	8 ; 64
47	0 ; 0	0 ; 0	120 ; 400	64 ; 400
48	- ; -	- ; -	512 ; 32768	256 ; 32768
49	- ; -	- ; -	- ; -	1 ; 1
50	- ; -	- ; -	16 ; up to 100 maybe	- ; -
52	- ; -	- ; -	80 ; 80	- ; -
53	- ; -	- ; -	16 ; 0	1 ; 4
54	64 ; 1024	- ; -	- ; -	- ; -
55	- ; -	8 ; 32	- ; -	8 ; 32
56	256 ; 512	- ; -	- ; -	- ; -
57	- ; -	128 ; 640	- ; -	- ; -
58	- ; -	- ; -	- ; -	4 ; 0
59	- ; -	- ; -	- ; -	500 ; 500
60	- ; -	- ; -	1-16 ; 96	- ; -
61	- ; -	- ; -	0 ; 0	- ; -
62	- ; -	0 ; 0	- ; -	- ; -
63	- ; -	256 ; 0	- ; -	- ; -
64	- ; -	- ; -	- ; -	16 ; 16
65	- ; -	- ; -	- ; -	0 ; 0
66	256 ; 1024	- ; -	256 ; 1024	- ; -
67	- ; -	0 ; 0	- ; -	- ; -
68	0 ; 0	- ; -	- ; -	- ; -
69	- ; -	- ; -	- ; -	1000 ; unlimited
70	8 ; 0	- ; -	- ; -	- ; -
71	112 ; 0	112 ; 0	- ; -	112 ; 0
72	- ; -	- ; -	- ; -	0 ; 0
73	- ; -	- ; -	- ; -	40 ; 40
74	- ; -	- ; -	0 ; 0	- ; -
75	- ; -	400 ; 2000	400 ; 2000	- ; -

Table 24: Ideal resource for the projects in 2010-2011.

project	CPUs	MemoryConfig	Memory (GB/processor)	Storage (TB)	Interconnect
1	thousands	distributed	8-32	1	default
2	32	shared	2	50	fast
3	0	distributed	0	0	default
4	500	distributed	0	0	default
5	400	shared	2-16	10	fast
6	24	distributed	2	1	default
7	200	distributed	4	1	default
9	128	distributed	4	4	default
10	4096	distributed	1	10	fast
11	10000	distributed	10	100	fast
12	1000	distributed	20	2	fast
13	10000	distributed	1	10	fast
14	64	distributed	8	1	default
15	64	shared	4	1	fast
16	200 000	shared	5-10	<1	default
17	500	distributed	4	20	default
18	200	distributed	1	1	default
19	12	distributed	16	4	fast
20	0	distributed	0	0	default
21	at least 500	distributed	32-64	1-10	fast
22	64	distributed	?	?	fast
23	2000	distributed	2	10	fast
24	5500	shared	4	100	fast
25	0	distributed	0	0	default
26	60000	distributed	unknown	unknown	default
27	256	distributed	8	10	default
28	256	shared	4	2	fast
29	16	distributed	1	1	default
30	16	shared	8	20	default
31	64	shared	0.2	2	fast
32	64	distributed	8	0	default
33	100	distributed	4	100 GB on home	default
34	256	shared	128	3	fast
35	200-300	distributed	0.5	2	fast
36	64	distributed	8	1	default
37	600	distributed	4	6	fast
38	512	distributed	2	1	fast
39	512	distributed	10	1	fast
40	64	distributed	2	10	fast
41	20	distributed	8	100	default
42	64	distributed	0	0	default
43	36	shared	16	2	default
44	8-16	distributed	32-64	1	default
45	30000	distributed	0	0	default
46	64	distributed	16	10	fast
47	2500	distributed	8	1	default
48	256	distributed	4	10	default
49	300	distributed	100	1	default
50	16	distributed	1-2	0.05	default
51	64	distributed	4	2	default
52	80	distributed	16	1	default
53	8	distributed	2	1	fast
54	640	distributed	8	5	fast
55	16	shared	2	100	default
56	128	distributed	4-8	1	fast
57	256	distributed	4	2	fast
58	4-634	shared	0	0	default
59	600.000	distributed	I am not sure	I am not sure	default
60	48	distributed	32	5	fast
61	128	distributed	2	2	default
62	20	distributed	2	1	default
63	0	distributed	0	0	default
64	4	distributed	64	1	default
65	0 10	distributed	0 1	0 4	default fast
66	512	shared	4	2	fast
67	8	distributed	0	0	default
68	0 20	distributed	0 2	0 1	default
69	1000	distributed	1	2	fast
70	0	distributed	0	0	default
71	0	distributed	1	10	fast
72	20000	distributed	?	?	default
73	80	distributed	2	1	fast
74	1	distributed	1	0	fast
75	2000	distributed	100	20	fast

Table 25: "Specify improvements that you believe are necessary as well as details about any 'Very Dissatisfied' answers given. Please specify the resource or site in case of any machine-specific remarks. Also specify software that you believe Notur should offer but currently is missing."

Number of comments: 20

- 1: Notur should offer open source software where available, because it would allow users to use the same software "at home" and on NOTUR resources. In my field, this would include Quantum-Espresso, SIESTA and GAMESS-US. The first two of these are for band structure calculations of crystalline solids, which is a strongly emerging field with importance for materials science.
- 2: We are using a mix of applications so "Application Performance" point 2 we have listed and ideal configuration for BCM since it will still be important for us in 2010. Storage is for permanent archiving.
- 3: The term processor, used above, is not defined! In our answer, we have taken it to mean "node". The diversity of applications used by our project makes it meaningless to single out an "ideal resource" in point 2 above.
- 7: Most of our projects are not the most demanding. Recently the 2-electron applications are reaching for us new level of requirements, but most of our applications fall very well within the limits of available resources.
- 17: Debugging has generally been a problem on Titan, as the Intel compilers provide very little feedback to the user. For instance, we also have access to the IBM xlf compiler at a different supercomputing center, and we have found it very useful to install our codes there, just in order to compile and test run with that compiler, as it catches quite a number of bugs that the Intel compiler doesn't. Further, it provides much more detail than the Intel one in terms of line numbers. Very useful.
- 24: The debugger GDB would be good to have.
- 25: No answers were given on 2. above. We typically run many jobs, each on 16 cpus. I'm not sure about memory requirements.
- 30: See before: storage of huge data-sets should be improved.
- 31: Our main requirement is a machine that can run at least 16 processors with shared memory (about 2 GB is required on that node), and where the rest of the system is easily using possible other nodes, utilizing MPI. Flexibility in the queue and run system is important for configuring the job.
- 34: It would be great to have some nodes on Njord with more than 16 processors.
- 46: Just a remark: I'm unsure what "Performance and debugging tools" and "General tools and utilities" refer to in the "Notur software resources" question.
- 47: There have been several problems with new version of the Intel compilers on Stallo, often not directly compatible with MPI versions, which has caused some significant problems in method development projects. These problems now appear to have been resolved. We also experience, due to the long run times of our parallel calculations, occasional problems are experienced in the message passing for running jobs, causing jobs to crash. This applies to all the machine we have been using, including hexagon and Njord. Our current usage (national+local quota) on Stallo corresponds to a continuous use of 1800 processors on this machine only. Not all application can utilize effectively a distributed system as given above.
- 48: Given that Titan and Stallo are technically quite similar, it is highly desirable to provide a unified user experience, especially to use the same queuing system on both machines so that jobs could easily be moved between machines.
- 49: The bottleneck for the applications I'm running is the cpu cache. Processors with larger and more efficient caches would allow us to reduce the cpu time per job.
- 53: I would very much appreciate a parallel version of Gaussian 09 on titan.
- 54: The GotoBlas library is efficient for our application. We prefer it to be installed on Hexagon.
- 55: Several versions of MPI libraries should be made available on titan.
- 66: In our molecular dynamics research, increasing the number of CPU cores is an efficient approach for increasing simulation size and its time-scale. However, the waiting time on the batch queues increases with increasing number of cores. Also we sometime have compiler compatibility problems, and software stability problems when we try to increase the number of cores for the MPI application that has in principle a great scalability up to more than 10k cores. In our case, 64 CPU cores on stallo was practical maximum. On hexagon, we didn't try so many cores because we had some compiler compatibility problems that we couldn't solve so far.

69: There is no debugger available on titan that knows about MPI.

73: We develop and test the code on a multicore desktop. This allows us to use titan exclusively for full scale runs. i.e runs of the engine where we are interested in the finally result. We develop, debug and test the code on a multicore desktop. This allows us to use titan exclusively for full scale runs. i.e., runs of the engine where we are interested in the final result. As a consequence we do not rely so much on debugging and performance tools on titan.

## 7 Notur User Support Services and Resource Allocation

This section of the questionnaire consisted of three parts. The first part included questions concerning satisfaction with the (basic/help-desk) user support services provided on each of the Notur hardware resources. The second part concerned advanced user support. The third part concerned satisfaction with the resource allocation and the established mechanisms for this.

### 7.1 Notur User Support Services

The projects were asked to state their overall satisfaction with several aspects of the user support services provided on each of the Notur hardware resources the projects have used. The following keywords are used in this section:

Support_TechAdvice:	"Quality of support given by the technical staff"
Support_ResponseTime:	"Response time for queries and problems"
Support_ResponseUrgent:	"Response time for urgent requests (e.g. file restores, quota changes, etc.)"
Support_Followup:	"Follow-up to queries and reported problems"
Support_TimeToSolve:	"Overall time to resolve problems"
Support_Accounts:	"Management of project and user accounts (creation, passwords, etc.)"
Support_Pages:	"Quality of support pages and other end-user documentation"

The answers are summarized in Tables 26 and 27. The two tables show a high satisfaction with the quality of the support given by the technical staff and the response times and follow-up to requests and problems, but at the same time a relatively low satisfaction with the quality of the support pages and other end-user documentation. In Appendix D, more detailed results are presented for each of the keywords per resource. More observations can be found in Section 2. The numbers in Tables 26 and 27 have the same meaning as the numbers in Tables 16 and 17.

Table 26: Satisfaction with user support services (1)

KEY	hexagon	njord	stallo	titan	AVE
Support_TechAdvice	4.3	4.4	4.4	4.0	4.3
pct	21%	32%	37%	37%	-
Support_ResponseTime	4.2	4.4	4.2	3.6	4.1
pct	21%	33%	39%	37%	-
Support_ResponseUrgent	4.5	4.3	4.4	4.3	4.4
pct	17%	27%	25%	27%	-
Support_Followup	4.3	4.0	3.9	3.8	4.0
pct	15%	27%	35%	33%	-
Support_TimeToSolve	4.2	4.2	4.0	3.7	4.0
pct	19%	31%	33%	35%	-
Support_Accounts	4.4	4.3	4.0	4.0	4.1
pct	21%	31%	32%	35%	-
Support_Pages	3.9	3.6	3.5	3.3	3.5
pct	21%	29%	33%	36%	-

Table 27: Satisfaction with support services (2)

KEY	VERYSAT	SAT	NEUT	DIS	VERYDIS	TOTAL	AVE
Support_TechAdvice	38	48	11	2	0	99	4.2
pct	38%	48%	11%	2%	0%	100%	-
Support_ResponseTime	34	43	22	2	0	101	4.1
pct	34%	43%	22%	2%	0%	100%	-
Support_ResponseUrgent	38	26	11	0	0	75	4.4
pct	51%	35%	15%	0%	0%	100%	-
Support_Followup	21	44	16	2	2	85	3.9
pct	25%	52%	19%	2%	2%	100%	-
Support_TimeToSolve	19	53	17	2	0	91	4.0
pct	21%	58%	19%	2%	0%	100%	-
Support_Accounts	39	34	12	6	1	92	4.1
pct	42%	37%	13%	7%	1%	100%	-
Support_Pages	10	38	37	5	3	93	3.5
pct	11%	41%	40%	5%	3%	100%	-

## 7.2 Advanced User Support

The projects were asked about their experiences with the advanced user support mechanisms provided by Notur. The answers are shown in Tables 28 and 29. Comments provided by the projects regarding organizing and improvements related to advanced user support are given in Table 30.

Table 28: "Are you familiar with the advanced user support services provided by Notur?"

KEY	YES USED	YES NOT USED	NO	TOTAL
Support_AdvancedFamiliar	21	26	28	75
pct	28%	35%	37%	100%

Table 29: "In case you have used it, were you satisfied with the results?"

KEY	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
Support_AdvancedSatisfaction	5	8	7	1	1	53	75	3.7
pct	7%	11%	9%	1%	1%	71%	100%	-

Table 30: "Please specify other services that advanced user support should include that will suit your needs in the best possible way."

Number of comments: 6

- 6: In our opinion, the advanced user support system in place already is close to optimal.
- 12: Optimization and scalability for large numbers (> 1000) of CPU's.
- 16: Introduction to advanced mathematical methods, computing points and weights of spectral methods. Advanced user support for new beginners on the supercomputers to get started.
- 24: Longer term support is better. The users will certainly appreciate collaborations with experienced scientific programmers.
- 59: We are happy with the way it is now scheduled, however, longer term, multi-year projects will make it a lot more feasible to plan a project and think further what can be done in the future. At CEES, we are in this project working with other disciplines as economists, statisticians and oceanographers. The questions and modelling-needs are continuously developing, thanks to output from Titan. It would be much easier to plan the future if we knew the support would be available over a long time-period.
- 75: Most of the work on the computers is done by PhD-students. They are in different phases, some newcomers and some well established. That means there is a need for continuous support and transfer of competence to and within the group. We are in great need for advanced user support which lasts over a longer period (years). We have good experience with involvement in the existing (months) advanced user support projects, and have a list of tasks still undone. Paying focus on the application side means we have a lack of experience and possibility to adjust to and follow the changing hw-development, risking to end in a deadlock with our sw.

### 7.3 Resource Allocation

The projects were asked to report their overall satisfaction with the mechanisms for resource allocation. The results are shown in Tables 31-32. The following keywords are used in this section:

- Resalloc\_Application\_Procedures: "Application procedures and feedback on applications"
- Resalloc\_Allocation\_Periods: "Frequency and duration of allocation periods (usually 6 months)"
- Resalloc\_Information: "Available information and instructions"
- Resalloc\_Amount\_Start\_Resources: "Amount of resources granted at the start of a new allocation period"
- Resalloc\_Amount\_Extra\_Resources: "Amount of extra resources granted (during allocation periods)"
- Resalloc\_Response\_Time: "Response time for applications for extra resources"

The responses are shown in Table 31.

Table 31: Satisfaction with Resource Allocation

KEY	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
Resalloc_Application_Procedures	12	43	10	4	0	6	75	3.9
pct	16%	57%	13%	5%	0%	8%	100%	-
Resalloc_Allocation_Periods	7	40	19	3	0	6	75	3.7
pct	9%	53%	25%	4%	0%	8%	100%	-
Resalloc_Information	8	47	9	4	0	7	75	3.9
pct	11%	63%	12%	5%	0%	9%	100%	-
Resalloc_Amount_Start_Resources	9	35	19	6	0	6	75	3.7
pct	12%	47%	25%	8%	0%	8%	100%	-
Resalloc_Amount_Extra_Resources	15	33	7	2	0	18	75	4.1
pct	20%	44%	9%	3%	0%	24%	100%	-
Resalloc_Response_Time	22	19	13	1	1	19	75	4.1
pct	29%	25%	17%	1%	1%	25%	100%	-

Table 32: "Specify improvements that you believe are necessary as well as details about any 'Very Dissatisfied' answers given. Please specify the site in case of any machine- or site-specific remarks."

Number of comments: 17

- 1: Keep up the good work
- 6: We are very happy with this aspect of the NOTUR system.
- 8: Support may be a little too technical, also the documentation.
- 10: The answers above based on the main application round this spring. The last round (allocation 10.09-03.10) have not yet produced any feedback or even allocation letter which is extremely dissatisfactory but given the good history, this last period is regarded as a one-time glitch in the system.
- 23: Our projects are of more long-term character (also from the funding periods, our projects last several years). Therefore, applications for 2-year periods would be more practical for us. However, we can live with the present practice.
- 24: NOTUR has worked very well in our perspective. There may have been some confusion for the large accounts that have a large allocation on one machine and a small allocation on another machine.
- 33: 12 month allocation periods (as recently adopted) are better for us.
- 34: Most of our time on Njord is from local allocation. For the period we are now in, we still have not gotten the final answer for our allocation application.
- 40: Shorter allocation periods, e.g. 3 months. Response time for allocations and information on allocation procedures should be improved
- 47: Whereas the application for additional computing time on Njord and Hexagon was dealt with swiftly, I have not yet been informed about the allocation for our project on the running resources, Titan and Stallo for the period 2009.2. Considering that we are now half-way through the allocation period without being notified, this is not satisfactory. Furthermore, while waiting for the final allocation to be made, we have been given 1/3 of our previous quota.
- 51: For the period (Oct 09 - March 10), the resource allocation committee plans to grant us the allocation units to the different HPC systems in Norway other than njord. My project was granted only 1/3 of the allocation units of previous period (2009.1 1 April - 30 September). Now, I have used up the allocation units on njord for period (Oct 09 - March 10). And, you have not given me any notice about the new HPC system I can use in the future.

- 55: Allocation periods should ideally be one year instead of six months. Because of the nature of the work and other unforeseen circumstances, granted resources might be underutilized for several months. With a longer allocation period such situations can be rectified with more vigorous use of the resources over the remaining period.
- 58: Re. allocation periods: This is a small project running for several years with only minor changes in resources. 6 month periods are too short. Recently, we had a one-year allocation period and this is probably the shortest allocation period being reasonable for a project like this.
- 66: Some information concerning to Notur resources application is provided only in Norwegian language. This might give a bit trouble for researchers from abroad like me. There seems some confusing description in English information, though the people in Notur are really helpful.
- 69: We were given the small amount of time we were granted for a new project in the previous period. So far titan hasn't been counting CPU hours, so it hasn't mattered yet. If titan starts counting CPU hours we will quickly use up our quota.
- 72: As someone who is new to these kind of computing resources parts of the application form were unclear. For example, I had never heard of tape storage before:) Also, it was unclear what we could apply for exactly. As for allocation periods, I think 6 months is definitely too short. We haven't even got ours up and running properly and four months have already passed. It is not quite clear to me whether we have to put in a completely new application after these 6 months. I would also like to add that it took me a while to figure out who I should contact for what kind of questions.
- 75: Need to plan long runs (several weeks/months)

## 8 Overall Satisfaction with Notur

In this section of the questionnaire, the projects were asked to state their overall satisfaction with several of the key components of the Notur project and how important they consider each of these components. The results are summarized in Tables 33 and 34. The following keywords are used in this section:

Overall_Hardware:	"Available computing hardware"
Overall_Storage:	"Available storage facilities"
Overall_Software:	"Available software"
Overall_Metacenter:	"User support organization"
Overall_Support:	"User support services"
Overall_Information:	"Available information and documentation"
Overall_Web:	"Notur website (www.notur.no)"
Overall_TechWeb:	"Technical support pages (docs.notur.no)"
Overall_Allocation:	"Resource (CPU-hour) allocations"
Overall_Notur:	"Overall satisfaction with Notur"

Table 33: Overall satisfaction

KEY	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
-	5	4	3	2	1	-	-	-
Overall_Hardware	20	44	7	0	0	4	75	4.2
pct	27%	59%	9%	0%	0%	5%	100%	-
Overall_Storage	10	37	15	1	1	11	75	3.8
pct	13%	49%	20%	1%	1%	15%	100%	-
Overall_Software	8	40	13	2	0	12	75	3.9
pct	11%	53%	17%	3%	0%	16%	100%	-
Overall_Metacenter	15	39	7	3	0	11	75	4.0
pct	20%	52%	9%	4%	0%	15%	100%	-
Overall_Support	18	39	8	3	0	7	75	4.1
pct	24%	52%	11%	4%	0%	9%	100%	-
Overall_Information	7	23	24	12	0	9	75	3.4
pct	9%	31%	32%	16%	0%	12%	100%	-
Overall_Web	6	30	27	3	0	9	75	3.6
pct	8%	40%	36%	4%	0%	12%	100%	-
Overall_TechWeb	5	22	27	5	0	16	75	3.5
pct	7%	29%	36%	7%	0%	21%	100%	-
Overall_Allocation	14	42	12	3	0	4	75	3.9
pct	19%	56%	16%	4%	0%	5%	100%	-
Overall_Notur	20	46	4	1	0	4	75	4.2
pct	27%	61%	5%	1%	0%	5%	100%	-

A number of respondents report dissatisfaction with the available information and documentation and the web pages. These two are probably related since dissemination of information and documentation usually takes place via web pages. The low marks correspond with the relatively low marks for the quality of the support pages and other end-user documentation shown in Tables 26 and 27.

Table 34: Overall importance

KEY	VERYIMP	IMPORTANT	NOTIMP	NA	TOTAL	AVE
-	3	2	1	-	-	-
Overall_Hardware	64	5	0	6	75	2.9
pct	85%	7%	0%	8%	100%	-
Overall_Storage	24	30	9	12	75	2.2
pct	32%	40%	12%	16%	100%	-
Overall_Software	23	28	14	10	75	2.1
pct	31%	37%	19%	13%	100%	-
Overall_Metacenter	25	36	3	11	75	2.3
pct	33%	48%	4%	15%	100%	-
Overall_Support	32	30	2	11	75	2.5
pct	43%	40%	3%	15%	100%	-
Overall_Information	24	42	1	8	75	2.3
pct	32%	56%	1%	11%	100%	-
Overall_Web	7	47	7	14	75	2.0
pct	9%	63%	9%	19%	100%	-
Overall_TechWeb	19	37	4	15	75	2.2
pct	25%	49%	5%	20%	100%	-
Overall_Allocation	55	14	0	6	75	2.8
pct	73%	19%	0%	8%	100%	-
Overall_Notur	48	14	0	13	75	2.8
pct	64%	19%	0%	17%	100%	-

Not surprisingly, the computing hardware is considered very important by a large majority of the respondents. Also resource (CPU-hour) allocation is considered very important by most respondents. More observations can be found in Section 2.

Table 35: "Specify improvements that you believe are necessary as well as details about any 'Very Dissatisfied' answers given. Please specify the site in case of any machine- or site-specific remarks."

Number of comments: 6

- 14: Impossible to find specific pages/information on the web site.
- 24: The centralized documentation page docs.notur.no is a good initiative and needs to be followed up. The advantage is that it is a search-able resource, but all the teams should play the game and put their documentation in there.
- 46: Despite a generous allocation for advanced user support, we were unable to utilize our given number of man-hours. It seems too few chunks of contiguous time were available for handling challenging tasks that take some time getting into.
- 62: It would be nice if you all can provide an initial guidance on how to start using a Notur facility by a one-day seminar.
- 68: Because of the delays in data collection we have had little contact with Notur.
- 73: We very much appreciate the completeness of the python environment (python, numpy, scipy and pypar were all available) on titan. This allowed us to get the code running on titan within hours. We very much appreciate the completeness of the python environment (python, numpy, scipy and pypar are all available) on titan. This allowed us to get the code running on titan within hours.

## 9 Summary of Experiences

The projects were asked to summarize their main experiences and satisfaction with the Notur project and comparison with other similar projects. The section consisted of three questions. The results are summarized in Tables 36-38. More observations can be found in Section 2.

Table 36: Question 1. "What does Notur do well? What aspects of Notur are you most pleased with?"

Number of comments: 53

- 1: NOTUR represents an important national scientific infrastructure. The idea of organizing HPC on a national scale is brilliant and has served us well for more than 20 years. The existence of NOTUR enables users to concentrate on science, since NOTUR handles many important technical challenges.
- 2: Provide access to high performance computing platforms that would otherwise be impossible for stand alone academic institutions. The maintenance and support is also much more professional than what the academic institutions could provide themselves.
- 3: Efficient use of resources wrt getting the maximum number-crunching capacity for the money.
- 5: NOTUR gives us an international advantage, since the application procedure is relatively simple relative to the size and importance of the resources. It is very important to keep bureaucracy low.
- 6: It is easy to access NOTUR to obtain help. The NOTUR personnel are keen to help and show a genuine interest in our problems.
- 7: The availability of the resources. Willingness to help. Positive attitude to the users.
- 10: Notur does well in providing stable hardware resources.
- 11: Relatively easy access to computing facilities and time.
- 12: User support and organization.
- 13: Availability of computational resources. Light administration.
- 14: Providing computational resources.
- 15: Most satisfied with the quick and precise technical support and the amount of high-quality application software accessible to the users.
- 16: CPU allocation procedures and machine stability.
- 17: For us, Titan is an absolutely essential part of our work; the analysis of QUIET data would simply not be possible for our case without these resources. Titan provides computing power with several hundreds of cores, available at any time. Further, the user support is absolutely outstanding, with an extremely helpful staff.
- 18: Available resources and overall reliability.
- 20: So far, we have only been using OpenMP type of parallelization and have been able to run on only 1 node. We are however looking forward to start using MPI with the help of the "advance user support". With the limitations that are imposed by our use of OpenMP, we are satisfied with Notur in general.
- 22: Support and quick answers
- 23: Problems are quickly fixed. The staff is responsive and helpful and also involve users in more long-term improvements of the computing infrastructure through consultations. Overall we feel that we are very well taken care of. Contact with NOTUR-members is always pleasant. They are no bureaucrats, but want that we as users/scientists have success. This is very nice.
- 24: We are most pleased with the quality of the hardware and the support services. The application for quotas and extra allocations is also working well.
- 26: In summary, we are very satisfied with Notur in general. The only problem we have encountered is that the waiting time at Njord can sometimes be a bit long.
- 27: Nothing special to report.
- 28: Computational resources and service.

- 29: Support
- 30: Good support, stable system.
- 33: Notur provides excellent resources and very good user support. The application procedure is straightforward and relatively efficient. Notur is flexible and adapts to our needs as they develop and change during the life of a project.
- 34: You have very good support and response to questions on allocations etc. The application for CPU time is straightforward.
- 35: Great computing facilities.
- 36: That I actually get computation time. It seems easier to get adequate computation time through Notur than to get money to buy a "home" cluster. I like that Notur has the possibility of applying for computing time for small projects also, not just large projects.
- 37: Good resources, easy access for cluster computing.
- 39: Using titan instead of njord (due to extremely long queueing time) has been crucial to be able to complete the project obligations.
- 43: Notur provides excellent supercomputing facilities. The software I use, can be installed and run on the platforms of NOTUR. Furthermore, I am getting support from staff of the supercomputers.
- 44: The daily operation of the compute resources is professional and very satisfactory.
- 45: Good and fast support, easy to get more cpu hours fast.
- 46: User support is outstanding. There is also a strong overall feeling that Notur wants to help us do our work, and I'm very grateful for it.
- 47: Notur provides ample computing resources of a diverse nature at a very affordable price (free ;-)). It is also clearly beneficial to us that Notur provides access to commercial software packages and not charging users for the cost of acquiring these software resources. In general, queue times are very good, and job turnover excellent. The structure for gaining access and computer time is rather informal.
- 48: Notur appears in general unbureaucratic and easy to use. The effort required to apply for small projects is reasonable. User support is generally good, we are particularly pleased with the help we got from the people in Tromsø. We are happy to see that it is now possible to apply for more than six months at a time.
- 49: Notur provides access to a large, high performance computing cluster. Access to the computing resources is reliable and well organized. Access to the computing resources of Notur has had been significant for the quality of my research.
- 50: Notur does well: To make sure that the system is running on a daily basis. Customer support.
- 51: (1) The system is running reliably. (2) Good technical support.
- 52: It was easy to get our code to work.
- 54: The performance of the supercomputers. The efficiency of the support.
- 55: Notur provides an efficient HP computing environment. Availability of machine time, speed etc are good aspects of Notur.
- 56: I am in general very pleased with NOTUR.
- 59: Notur is doing very well, they are good in information, professional, positive and flexible.
- 60: Support and available hardware
- 62: The facilities.
- 63: Very good technical support. Generous quotas.
- 66: The system is generally ok. Notur fits well to our requirements of parallel computing with MPI.
- 67: Easy to obtain resources allocation to students and post docs.
- 72: I am pleased with the availability of the resources Notur offers. With the current revolution in genomic data it is essential for researcher to have access to such resources.
- 73: We were most pleased with the ease of getting access to titan and the generous allocation of resources. I can not over emphasize how important NOTUR/titan has turned out to be for our work
- 74: So far I have not used the computer and Notur much - therefore many unanswered questions
- 75: Stable resources- Access to well experienced support people

Table 37: Question 2. "What should Notur do differently? How can Notur improve? Are there specific issues where Notur can do better (services, support, etc.)?"

Number of comments: 39

- 1: User support is rather unpredictable, and relies on personal connections to the people that carry out the actual work. In my experience the support is very good once I manage to sell my problem to a competent person. It is hard to give specific advice here, because the need for a public ("free") service in principle is unlimited.
- 2: Our projects requires large amounts of permanent archiving storage. We hope NorStore will eventually provide a good solution these needs in the future. The establishment of NorStore has taken more time than indicated and we still do not have a clear picture of what the future capacity will be. A clear plan for this would be beneficial for us.
- 3: It may be considered to delocalize parts of the package-specific user support (and package installation and maintenance work). This means that service personnel should get programs/packages allocated as their general responsibility, that is, the responsibility should extend to machines outside of the local computer center. In this way, they may develop specialized competence wrt user support, and be able to offer better support especially wrt run-time conditions/parameters. Up-to-date documentation would be a good idea. It is generally not the case now. The web-pages (technical) are very often outdated and lacking in detail.
- 6: Compared to high performance computing centers elsewhere in the world, NOTUR does very well. Within our research group, we see no need for changes in the NOTUR organization.
- 7: Overall, there are no single issues. Perhaps the changes in procedures, changes in deadlines should be avoided. The information about changes and deadlines could be broadcasted in good time.
- 8: Better (less technical, but still informative) and more frequent info from Titan would be welcome.
- 10: User support can be improved but this is not crucial for our group. The same can be said for advanced user support. In a scenario with limited funds it is important that the majority of available money is used for hardware.
- 11: User support in particular performance profiling and optimization can be given more emphasis. Larger systems are dearly wanted!
- 12: As always, more CPU resources are welcome.
- 14: More structured and intuitive web portal.
- 15: Have no suggestions.
- 16: Higher number of single processors available for code development.
- 17: Improved stability is the number one issue for us. Smaller points are improved debugging tools, resolution of the long-standing MPI problems we see on Titan, and a better queue system, in which the owners of nodes actually get access to the respective cores, and won't have to stand in queue together with non-owners. This doesn't seem to work optimally.
- 18: I would have preferred that there was more flexibility regarding when applications for computer time can be submitted. I guess I sometimes would have liked more "interactive" support. When setting up new software, it is sometimes not a single support request to get something fixed, but a bit of back-and-forward to get things working. In these cases it could be useful to have a bit more time with support.
- 23: It is difficult to find the actual deadlines for application for CPU times on the web. The changes of the application procedures can cause some confusion.
- 24: The documentation pages could be improved.
- 26: Improve the capacity at Njord.
- 27: Technical documentation should be improved.
- 28: Documentation and website.
- 29: Batch waiting time, but I guess this is hard to improve without buying new hardware.
- 30: Please add more storage / a bigger archive system.
- 33: We prefer to apply for CPU hour resources once per year than every 6 months.
- 37: Visualization facilities would be nice to have access to

- 39: Can possibly improve documentation on Internet.
- 43: Right now I do not have such points.
- 44: Again, we are very happy with the daily operation for the data centers. However, the application and reporting procedures could maybe be simplified (the new web-based application form is a major step in the right direction).
- 46: Disk IO has been a major bottleneck, and Titan has periodically been struggling with interruptions. However, I'm generally satisfied with their efforts to solve those problems. It is crucial that information on ongoing problems is put on hpc.uio.no and the login screens as soon as possible, so we can work around any interruptions; this also avoids duplicate error reports.
- 47: Information flow could perhaps be better. Although it is good that e-mail communication is limited, I have the impression some important information is not provided in a manner that is easy to take note of (for those of us who do not visit the Notur web pages on a regular basis). Requires that the user is more aware of the normal deadlines for applications.
- 48: Documentation is obviously never good enough ;). Would there be any way to draw more on the community in creating documentation? Especially, to ensure that questions individual users had, and answers to them, become available to all?
- 49: Feedback on applications and quota allocations could be better. There is also some room for improvement regarding new accounts and users.
- 50: The internet pages are a little bit messy. But I rarely use them.
- 51: (1) I hope to have some interactive node to code debug and development. (2) Notur should organize some workshops on how to use the systems more efficiently.
- 54: The parallel efficiency of Stallo for our application (the VASP code) is not very high? We think it is due to the narrow band of the Xeon 54XX series processors. Is it possible to update the processors to the newer Xeon 55XX processors?
- 56: No comments
- 62: Some extra documentation of machine architecture.
- 66: System information of the Noture computing resources appears distributed over the internet, reflecting probably the real locations over the campuses/organizations. From the users' side, this information could be unified into a Notur web-site, then we could find faster/easier a best solution for a specific computing problem research.
- 72: For me part of the problem is that I am so new to these kind of computing facilities, and such use of computing power in general. I would not be surprised if more researchers encounter this problem, or maybe they just avoid using the Notur resources. In that respect I feel that such newbies as myself are not part of Notur's target audience (yet), and for myself it has thus been quite difficult to find my way around the resources that are offered and to understand how I can use them best.
- 73: I am overall very satisfied with NOTUR/titan. I would be hard pressed to find an area apart from my little comment about the queuing system, where NOTUR could do better, than they already do.
- 75: Be in contact with the users in order to understand the priorities the users have. That will help being flexible when special needs occur (special testing phases or needs where the queue system is obsolete).

Table 38: Question 3. "How does Notur compare to other centers that you have experience with? What do other centers provide that Notur should provide as well? Please list the center(s) that you use in the comparison."

Number of comments: 30

- 1: I have no relevant experience here.
- 2: I have no similar centre to compare with.
- 6: See answer above.
- 7: We do not have much experience there. For years now, Notur has been both sufficient and essential for our activities.
- 10: Notur does very well. Centers for the comparison: NASA-AMES (problems with stability and terrible problems with security issues).

- 14: I have experience from Sweden (-1994) and Denmark (-2002). Notur compares very well. In both Lund (Sweden) and Denmark, we had specialist scientists (e.g. in computational chemistry) hired by the computer center to do advanced user support/research.
- 15: We have no experience with other HPC centers.
- 17: I've only got experience with NERSC as an alternative, which is a bigger and more "professional" center. From my experience, it is much easier to get things done on Titan, in terms of user support and updates. On the other hand, stability has been much worse on Titan.
- 18: Have used a center at the University of Florida, and Notur is in most respects better.
- 23: I think that NOTUR does overall an excellent job. Years ago, I worked with DKRZ computers in Germany. Also the DKTZ did a good job. I have no spontaneous ideas for improvements. In the long-term, as computing gets more and more complex, we have to think in a concrete action on how to use the evolving complex systems more efficiently. Remember, that we often hire post-docs or PhD students, who then have to work on a new system. It would be great to train them better for the machine usage (including also myself!), or to facilitate the use by automatic software systems (reliable autotasking compilers or the like).
- 24: We have not used other centers recently. Our users coming from China, Australia and Canada seem to be quite happy with NOTUR.
- 27: Do not have experience with other comparable centers.
- 28: Comparable with the HP-center at Imperial College London.
- 30: Comparison with the DKRZ (Deutsches Klimarechenzentrum) in Hamburg, Germany: they provide more or less unlimited storage capabilities for data.
- 33: This is the first center I have been involved with but my feeling from colleagues is that the availability and allocation of CPU resources is more efficient and flexible than many computing centers. That means we can use the resources themselves more efficiently.
- 37: It would be nice to have a visitors' program, where users can come and sit in a space adjacent to, or amongst, the consultants, so that intense debugging and porting can take place interactively and quickly. Each of the centers (Oslo, Tromsø, Trondheim, and Bergen) should have such a space, which could be booked for up to a weeks' time. In return for the consultants services, the user/visitor should then give a presentation to the consultants about their work. I've experienced this at ARSC (Fairbanks), NCSA (Champaign/Urbana) and SDSC (San Diego).
- 43: I am satisfied with the service and facilities provided by NOTUR.
- 44: No experience with other centers.
- 47: Compared to the other international supercomputer centres that I have had access to (NSC 1997, CSCS 1997, SDSC (1999-2000)), Notur appears much more flexible in providing access and resources to the users. This probably is in part due to the Norwegian supercomputing system, which has the dual role of providing capacity resources for all supercomputing activities, as well as capability access, whereas other centres have focussed on providing the high-end resources, leaving regular use to each user to find local solutions. For us as scientists, this allows us to focus on science rather than operating local computer systems.
- 48: We do not have direct experience with other centers, but from discussions we have had with colleagues abroad we have gained the impression that NOTUR is easy to use: reasonable amounts of paperwork and direct access to relevant and eager support personnel.
- 49: My project co-workers have experience with the computing resources at the universities of Leeds and Leicester in the UK. Compared to those, Notur appears to have significantly less down-time, and a queue system which is more transparent and predictable.
- 50: I don't know.
- 52: We have only used Notur
- 54: We only benefit from the Notur project.
- 56: My impression is that NOTUR is very well organised.
- 62: Ok, but not the best. I have only worked at Supercomputer Education Research Center, IISc, India. There the facilities, software and helping personnel are better than that we have in Notur. (<http://www.serc.iisc.ernet.in/>)
- 66: I have some experience with CSC (Finland) and some computer centres in Japan, and the Notur system is very much OK in general with their high performance computing hardware. Comparing with CSC that keeps all their HPC systems in under control in their own building/facilities, the Notur HPC resources are distributed over the universities in Norway. This is pretty much understandable for research-driven developments, but from users' side, it could have unified information on the systems.

72: I have no experience with other centers.

73: I just thought I use this space to share my experiences with queuing systems that have a suspend and resume policy in place. Before coming to Norway I was running some small jobs on clusters belonging to ANU/APAC and their clusters had a suspend and resume policy in place. That would mean that jobs like the ones I run would be suspended by the queuing system if somebody was asking for resources for a short job or an interactive login and when the small/short things were finished the large/long jobs would continue. For me it really does not matter if I get my results in 7 or 9 days. I kind of liked their approach very much and it made it easier for large long jobs to coexist with small short jobs in the same queue. [http://anusf.anu.edu.au/dbs900/PBS/local\\_modifications.html#preempt](http://anusf.anu.edu.au/dbs900/PBS/local_modifications.html#preempt)

75: Very well. Compared to colleagues they are often larger groups and have (partly) access to own hw, but most important do they often have a wider competence (computational science-like) to take care of the needs coupled to the application itself (data transfer, parallel issues, efficient storing, visualization,...). We have to establish separate contacts with different groups in order to follow up - This is time consuming and not efficient and results in less knowledge assimilated in our group. In the long run, we loose!

## 10 Future Needs

The projects were asked to specify any known or expected change in needs for resources within their project in the coming 24 months. An answer was not requested for projects that will cease to exist in 2009 or will have similar needs as today. The section consisted of six questions. The answers are summarized in Tables 39-43. More observations are included in Section 2.

Table 39: Question 1. "Please give a brief description of any known changes in your research activity in the coming 24 months that will affect your needs for computational and storage resources and user support. This could for example be the termination of projects, start of new projects, specific deadlines, etc."

Number of comments: 55

- 1: I envision business as usual for at least the next 24 months. Students come and go, but their number is steady at between 1 and 3. When they are few, I have more time to do calculations myself, so the variation is not so big.
- 2: No big changes are anticipated.
- 3: Expect man power at present level.
- 5: The project portfolio is steadily increasing, and we expect our needs to increase accordingly.
- 6: We have problems that are based on repeated inversion of large sparse matrices. Any means to improve the performance for this class of problems would be very welcome.
- 7: As mentioned, the new 2-electron approaches are much more demanding than previous applications. They will mean an extension of our computational needs. It is too soon to assess at present to which depth we will follow these studies.
- 8: No big changes, although the memory and CPU requirements will increase as the model resolution increases.
- 10: The group has expanded with one faculty position and the development of our radiation MHD code has resulted in a production quality code. Results so far has put the group as a central modelling group for a future space mission, IRIS, to be launched December 2012. The need for computational resources will therefore increase and we can basically use whatever we can get.
- 11: We'll continuously attempt to solve larger and larger problems, using more and more processor cores. Many-core (at least 32 cores per node) and GPU-based architectures will be of great research interest. Therefore, larger systems are wanted!
- 12: Application of a new type of ocean model for testing before putting it in operational use.
- 13: We are in a process of increasing the resolution of our model. The extend to which the improved model can be used, will be limited by the available computational resources.

- 14: The needs vary a lot depending on the amount of students/postdocs and the phase of the various subprojects.
- 15: In the coming year we will use less commercial and public domain software, and primary apply our own computer models.
- 16: Similar needs as today.
- 17: We are about to scale up the QUIET experiment from hundreds to thousands of detectors, as we move to Phase II of the project. This will eventually require an expansion of our disk space from tens to hundred TB. However, this will only happen in two or three years, and hopefully we will have established a working solution by then.
- 18: I expect somewhat higher level of activity (large number of single processor jobs) but no great changes.
- 20: Our need for computational resources will increase. We are more people (5 persons) using the nnXXXXk account and with use of MPI, we will be able to run simulations with higher spatial resolutions. Two of our projects will terminate in 2010 and two in 2011. We expect to make several climatic runs for all projects.
- 22: 5-7 people will work with Notur facilities in the next 24 months. A new project starts up in February 2010. Larger molecular systems will be studied.
- 23: We expect that our models will increase in resolution and complexity (more and more Earth system components will be coupled together). We expect more ensemble runs (identical runs with slightly changed initial conditions) and also over longer periods. Further we anticipate that data assimilation of observations into coupled climate and Earth system models will become important. All this means a dramatic increase in CPU time, storage demands, processor speed and memory. We could sustain our funding in the project for the next 24 months even with the termination of the EU FP6 Integrated Project CARBOOCEAN which we coordinated.
- 24: The project will continue over the next 2 years based on existing projects. A few proposals are in the pipe. The needs are expected to continue growing, say, linearly.
- 26: No known changes.
- 27: Will most likely have increased need for resources due to increased activity in High Throughput Sequencing. This is discussed with Notur/Norstore through the FUGE Bioinformatics platform.
- 28: We will have similar needs as today, but we expect a need for more computing time.
- 29: I need to complete the calculations for my PhD within this semester, but I plan to continue a collaboration with the Department, so my needs would be unchanged.
- 31: The project ends in 2010, and we need a lot of CPU time to finalize our work. It is still unclear if we can manage to complete all simulations. This is caused both by technical problems with the model system, and too small allocations that forces us to minimize both simulation lengths and tests to a minimum. Hopefully new projects will come where we have the possibility to upgrade the model system to utilize a capacity machine.
- 33: We hope to have 2-3 new postdocs working on the project during 2010-2011. This will allow us to expand our project and increase the resources we require.
- 34: We currently have a cooperation with Georgia Tech in an NFR funded project. In this project, most of our distributed memory CPU time is used. That is, all the time on Hexagon, and a lot of our time on Njord. The NFR project ends next year. We are not certain what will happen after this. Hopefully, we will get new projects, but there is also a possibility that our need for distributed memory CPU resources will decrease.
- 35: End of project is January 2011.
- 36: I will continue to use Quantum Monte Carlo codes to study various problems. Beside this there is a chance that I will start using more memory-intensive applications in the next 24 months as I have plans of implementing a variational approach which essentially is a minimization algorithm with very many degrees of freedom.
- 37: While the mix of simulations I do will change from impacts and tsunamis towards explosive volcanism, there should be no substantial change in the resources required, either in terms of cpu, memory, or storage.
- 38: We finally managed to get PhD positions for projects that will use the codes that we developed. Previous projects mostly focused on the development and could not make enough use of the developed tools. This is/was a resource problem on our side and will be improved now. As a result we should be able to run jobs more regularly and produce more science with the hours that we get.
- 39: My project ends in 2009.
- 41: No changes
- 43: In addition to my previous scientific activity related to study of novel materials for semiconductor electronics, defects in TCOs, another topic has been included. The topic is the study of defects in Silicon. The previous activity still is going on, since the project we have will be continued until 2013.

- 44: Basically, we adapt our workload to the available compute resources. Larger machines with more memory results in bigger experiments, more results and better research quality.
- 45: Our group will start a new analysis of MR-images in 2010, analysing the images of 1000 human brains. This will probably lead to an increased use of stallo in our case.
- 46: We have used a tiny fraction of our allocation so far. That's embarrassing; however, I see that we're not alone. In any case, we believe that will change in the next few months.
- 47: We will continue to expand somewhat, and thus we expect a somewhat further increased need for CPU cycles, but not a dramatic increase. We expect an increasing use of capability platforms due to a slight shift in the scientific profile of the center. Assuming the grid project devoted to Dalton will have a successful outcome, we may shift jobs to the NordGrid and NorduGrid platforms if the services provided are reliable and give improved run-time for our applications.
- 48: We expect to ramp up the amount of large network simulations we will perform, increasing your CPU and RAM requirements. We will likely be running larger numbers of jobs with O(64) cores, but may also go for even larger jobs.
- 49: My project receives financial support from the Norwegian Research Council until the end of 2010. During 2010, we will probably need additional computing time as we finalize ongoing projects.
- 50: No changes.
- 51: We just started a new three-year project (supported by Norwegian Research Council and industrial sponsors) on Advanced Computing of slug flow. Hence, we wish to get continuous support from Notur.
- 52: We will proceed as planned
- 54: We will continue the present project in the next 24 months.
- 55: Main concern is availability of large and long term (a few years) storage capacity.
- 56: My group will request regular use of 4-6 parallel 64-128 jobs on a regular basis. I have two big projects that require large scale parallel cfd. One ending end of 2010, one ending 2012.
- 59: We have to finish this project during spring 2010 as a PhD-student will submit her thesis by then. We had 70 hours of advanced user support for fall 2009, but we ended up using 105 hours. We have therefore applied for 200 hours in 2010 as this will need high computation effort in order to reach the deadline.
- 63: We had a peak activity in 2008/2009 with my sabbatical at MIT, and we are now gradually establishing atomistic- and multiscale material modeling activities at NTNU. At present we have no strong opinions, but look forward to the upgrading of Njord.
- 65: A large multi-user software system has been developed in my group during the last two years, and this system is now increasingly used internally and will soon become known also to external users. If successful, we expect that our system will soon generate a large computational load, which will also have peaks due to varying user activity. We also expect the increased use of the system to lead to increased storage demands. As we have a developing system, and have user-initiated runs, it is not trivial to make use of the Notur resources and thus user support is and will be very important.
- 66: We have started our molecular-dynamics-based research work 2009, and extend this during 2010. The fundamental parts of the task have are accomplished with several publications in this year. Next year we would explore larger simulation systems aiming to applications on the project with more computing resources.
- 68: In one month time we should have the data, we would like to transfer the allocation to the titan servers and proceed with the analysis.
- 69: One of the uses of our code involves running large 3d models (500 or 1000 cells per dimension). We have tested the code on smaller versions of the model and have got to the point of being ready for larger runs. We want to explore parameter space with 500 cube runs and then pick a small number of configurations for more detailed modelling.
- 72: We will have a great increase in data coming from next generation sequencing machines. This data will have to be stored, and analysed. Our project has only just started and will run for another 3.5 years.
- 73: The project we use titan for is expected end in January 2011. We do not anticipate big changes in the need for computational resources.
- 75: Need for more resources being used as efficiently as possible. This will be connected to available advanced user support.

Table 40: Question 2. "Will the need for computational resources change significantly in your project in the next 24 months? If yes, please state the reason(s) and detail the changes in your needs."

Number of comments: 48

- 1: In my field, computational resources remains a major bottleneck. To produce research at a high international level requires HPC infrastructure at a high international level. As long as NOTUR remains at such a level, so will our group :-)
- 2: The use of BCM will gradually be reduced and our demand for capability platforms will accordingly be reduced.
- 3: The need for computational resources is expected to increase at the same rate as during the past few years, as we attack problems of increasing size and complexity.
- 6: We foresee no drastic changes. We are striving for larger and larger systems as we are after the infinite-size asymptotics of our systems. This entails faster and faster machines and better and better codes.
- 7: As mentioned, the new 2-electron approaches are much more demanding than previous applications. They will mean an extension of our computational needs. It is too soon to assess at present to which depth we will follow these studies, but it seems that these demands will bring our project to a different class of needs.
- 8: See previous question.
- 10: See above. We need to run several runs with 1024-2048 cores over many months.
- 11: If possible, we'd like to use thousands of processor cores much more frequently. This is both for performance studies and production runnings.
- 12: The application of a more sophisticated ocean model will increase the requirements for more CPU power. In addition, a planned increase in the model resolution (finer grids) will also require more computational power.
- 13: Yes, more CPU hours will be needed.
- 15: As a consequence of our answer to the former question, our jobs will typically run on fewer processors and our I/O intensity will be much lower. Thus storage capacity will be less important. Our computer models involve large matrices, however, so available memory during computation will still be important.
- 16: Similar needs as today.
- 17: As far as computing power goes, we probably have enough with what we have already. Most likely, though, we will probably expand slightly.
- 18: No.
- 19: The amount of cputime needed will increase significantly - the size of the research group has tripled.
- 22: Increase due to more detailed model systems.
- 23: We participate in a series of EU and NFR projects on Earth system modelling (EU: COMBINE, EPOCA, COCOS, MONARCH-A(under negotiation), a CarboOcean follow-up (submitted)). We will participate with biogeochemical modules to the 5th IPCC assessment report. The already available funding for project work on Earth system models, feed back analyses, and future climate projections require a steady increase of computational resources.
- 24: We may become involved in assimilation activities with the Norwegian Earth System Model, in collaboration with the Bjerknes Center for Climate Research. These new applications should not change the nature of our needs, but they are expected to be quite demanding in CPU allocations.
- 27: See above.
- 28: No.
- 29: No
- 30: No.
- 31: Hopefully, we are going from a model tool that can run on a capacity computer.
- 33: Since our current work has been very successful, we will request more resources to run larger models and explore more conditions for our existing simulations.
- 35: no

- 36: I will still need computational resources, but my needs will probably not change much. However connected to the variational approach I might need more RAM, then also it might be necessary to use more temporary storage space on each node. Actually it might be favorable for this project to have fast local harddisks.
- 37: No real changes in needs.
- 38: Stay constant.
- 40: Our need for computational resources are gradually increasing as the simulation packages we utilize are getting more computationally demanding..
- 41: No changes
- 43: No change.
- 44: No
- 47: We will continue to expand somewhat, and thus we expect a somewhat further increased need for CPU cycles, but not a dramatic increase. We expect an increasing use of capability platforms due to a slight shift in the scientific profile of the center.
- 48: There may be a certain shift in jobs, from embarassingly parallel to MPI-parallel.
- 50: No changes.
- 52: We applied for more than we were allocated. We thus need more.
- 54: No.
- 55: No
- 56: Large 3D CFD simulations, see comments above.
- 58: No and yes. Our overall resource demands are small. We will soon install a new application on our own account at titan and if this works successfully it may alter our demands but probably only marginally. This will be considered in our new applications.
- 59: Yes, we will double the hours for advanced user support up to 200 hours and will also need accordingly increase in computation time.
- 60: No
- 65: Due to the above described evolution in the project, we expect a significant increase in the need for computational resources during the next year, and continuing into 2011.
- 66: Our simulation systems should be increased in the size in both space- and time-scale, and thus it will require more CPU cores. Now we use about the 64-128 cores, but may need to increase the number to 2-10 times more according to the simulation systems in the next year.
- 69: See answer to previous question. Larger models will use more CPU hours.
- 72: Yes, see above. We will require increased storage capacity and depending on the analyses possibly more CPU hours.
- 73: There might be slight increases in the need of computational resources depending on which changes/enhancements to our code we would like to try in the coming months in full scale runs on titan.
- 75: We are in the position to make use of whatever is provided, but also need backup of competence in fields that are tangent to efficient utilization of our application-sw.

Table 41: Question 3. "Will the need for storage change significantly in your project in the next 24 months? If yes, please state the reason(s) and detail the changes in your needs."

Number of comments: 49

- 1: No. By temporary measures, our needs are not especially large.
- 2: Probably not.
- 3: Yes. Temporary storage: see previous item. Permanent storage: considerably more, due to increasing demands (from central authorities) wrt long-term documentation of research results and in particular primary data.
- 5: We use today a temporary solution at the University of Oslo, but we will probably need a safe long-term storage solution within the next 24 months.

- 6: No.
- 7: This will increase, but compared with other projects probably quite moderately.
- 8: It will increase, although I cannot say by how much.
- 10: No change. Only need for temporary storage.
- 11: The storage need will increase accordingly, but not dramatically.
- 12: By increasing the grid size of the models, the demand for storage will increase, a factor of 4-10 should be expected.
- 13: No (we already use norstore).
- 14: No needs.
- 15: Yes, it will be reduced.
- 16: Similar needs as today.
- 17: See above.
- 18: No.
- 19: About the same
- 20: We need several TB, but so far we have not applied for storage resources.
- 22: Increase
- 23: Yes, due to the increasing production scenarios (once the present models under development will get more operational) we expect a significant increase in need for storage. We also will need long-term archiving facilities which are accessible by multi-users (at best through online retrieval tools such as CERA at the WDC Climate in Hamburg).
- 24: Same as for CPU time, we are expecting a linear growth with time.
- 27: See above.
- 28: No.
- 29: No
- 30: See remarks before: More storage would be nice.
- 33: We have norstore resources that are sufficient for the near future.
- 35: no
- 36: Not permanent storage.
- 37: No real changes in needs.
- 38: No
- 41: No
- 43: No change.
- 44: No
- 46: We believed we it would be useful to store vast amounts of intermediate computations from physiological simulations. However, with disk io being as slow as it is, we may have to rethink that.
- 47: No
- 48: We may need more storage, but believe that we will not enter the group of "large" storage users.
- 50: No changes.
- 51: No change.
- 52: No
- 54: No.
- 55: Yes. Large outputs from Earth System models will need large storage. With our project coming to an end in 2010, we will need increased flexibility in post-processing of model outputs data.
- 56: no
- 59: No, not really. The results will be downloaded for each large simulation.
- 60: No

- 65: We currently expect a somewhat increased need for storage, due to increased activity with our system and increased interactions from external researchers. We currently expect a somewhat increased need for storage, due to increased activity with our system and increased interactions with external research groups.
- 66: Our project produces much output data that should be kept at least until publications. So far we have about 1-2 TB data on stallo, and it is increasing.
- 69: See answer to previous question. Larger models generate more data (120 GB per model for  $500^3$ ). The data will require medium term storage (i.e., until the papers based on the results are accepted) for analysis.
- 72: Yes, see above.
- 75: We have recently started using Norstore, and through advanced user support project we expect to use it a lot. The next 24 months, 500TB.

Table 42: Question 4. "Will your project need access to hardware resources in 2009-2011 that are currently not provided by Notur (e.g., other computing architecture, other type of storage, etc.)? Please detail your needs."

Number of comments: 45

- 1: No. More of the same, please!
- 2: Probably not.
- 3: Not new architectures, but we do fear that the present pool of machines may become clogged towards 2011 and hence needs to be expanded in terms of node number and disk/memory. Moreover, as fatter nodes become shelfware, present clusters should be expanded with these (to some degree).
- 6: No.
- 7: Not essential, according to my opinion.
- 8: Not that I'm aware of.
- 10: Just more cores.
- 11: More processors and many-core and GPU architectures.
- 12: MPI platforms with a large number of CPUs, i.e. between 1000-10000 CPUs.
- 13: More CPU hours will be needed.
- 14: No
- 15: No.
- 16: No
- 17: Not applicable.
- 18: No
- 19: no
- 22: No
- 23: We cannot so often migrate our complex models to new architectures. Therefore, systems such as hexagon and njord are fine for MPI or OpenMP.
- 24: Not as of today's situation. But we expect Hexagon will become completely packed at some stage and then we will have a problem and more resources will be welcome.
- 26: No.
- 27: Not as far as we know.
- 28: No.
- 29: I hope to be able to access a linux cluster at NTNU. This should help running OpenFOAM (which is mainly designed for linux).
- 33: Not that we know of.

- 34: It would be good with a shared-memory system with more CPU's than what Njord has. This has to be a system with very fast buses, just like on Njord. A system like for example Stallo, with a very slow bus, will not be useful, even if more CPU's are provided for each node. This system scales very poorly in comparison with Njord.
- 35: no
- 36: No
- 37: No.
- 38: No
- 41: No
- 43: No, will not.
- 44: No
- 47: We observe the gradual incorporation of GPUs as computing platform. Having access to a test system for developing and testing new algorithms for these architectures, though we do not foresee a shift to this kind of compute platform in the near future.
- 48: No.
- 50: No.
- 51: No change.
- 52: No
- 54: Previously, we have applied for the resource Stallo. However, we found it does not fit our application for the aforementioned reasons. Therefore, Hexagon is our best choice now.
- 55: No
- 56: no
- 59: We have a self-developed program.
- 60: No
- 66: Molecular dynamics in principle is a good subject for parallel computing. The number of cores (and waiting time of the batch queues) is most important for us.
- 72: Not that I am currently aware of.
- 75: At present, more of what is already around. There is a potential in extending our knowledge in CFD, but then we need cooperation with resources from other fields (computer science-like). That may open possibilities for other/new architectures and storage.

Table 43: Question 5. "Will your project have a need for services in 2009-2011 that are currently not provided by Notur? Here, one could think of for example other types of support, other types of resource allocation, services for visualization, grid-enabling of applications, etc. Please detail your needs."

Number of comments: 38

- 1: See previous point.
- 2: Probably not.
- 3: There is need for much improved user support and shorter response times of these, rather than new types of support.
- 6: No.
- 7: No essential additions.
- 8: Probably not.
- 10: No
- 11: More involved performance profiling and optimization.
- 15: No.

- 16: No
- 17: Not applicable.
- 18: No
- 19: no
- 20: At XXXX, we run projects that are not supported by the Research Council of Norway (RCN). It has therefore been necessary to invest in our own small cluster of CPUs. These are typically applications that are supported by Aqua culture authorities, and could easily be run on NOTUR computers. It would have been beneficial to know NOTUR's policy on this subject. What is the cost of using NOTUR's resources for non RCN-projects?
- 22: ?
- 23: A link to the World Data Centre network would be good to link climate data and other geoscientific data sets (ocean, atmosphere, land, ice etc.) from model runs and observations with others.
- 24: Not as far as I can see. A central data server or a central visualization platform could be interesting. The grid-enabling of our applications could be good but we do not have projects to work on that as of today.
- 26: No.
- 27: Not as far as we know.
- 28: No.
- 33: Visualisation is very important for adequately displaying and communicating our resources. We have not spent time considering visualisation tools yet but we would definitely benefit from visualisation services.
- 35: no
- 37: Visualization facility would be nice.
- 38: No
- 41: No
- 43: I am going to apply for advanced user support. The reason is that I am only user and I feel difficulty in installation of new updates of the software, which I use. If I spend more time for installations and updating the script files, I can do the job. However, I prefer to spend more time for research work. For the software installation purposes I prefer to get support from qualified people of NOTUR, which know very well the different features of platforms.
- 44: No
- 48: No.
- 50: No.
- 51: No.
- 52: No
- 54: No.
- 55: No
- 56: no
- 59: No needs.
- 60: No
- 72: We might need some more support - I haven't applied for advanced user support yet but will do that in the near future. Our project would benefit from a shared project area to which all a number of users have access. This was currently not available (although that now being sorted). I can imagine that other projects would benefit from such shared area structures.
- 75: Yes. We are desperately (constant) in need for compatible competence. Daily we feel possibilities of doing new cases disappear because we dont have capabilities to take larger steps beyond the limits we feel we have control over today (computational wise). I.e. trying other Notur-resources, extend number of cpu's, verification/validation/performance on new hw, efficient storage possibilities,...

## A APPENDIX for Section 5

Table 44: [Hardware\_Uptime]: "Machine availability (uptime)"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	11	7	1	0	0	2	21	4.5
pct	52%	33%	5%	0%	0%	10%	100%	-
njord	14	9	3	0	0	3	29	4.4
pct	48%	31%	10%	0%	0%	10%	100%	-
stallo	19	15	0	1	0	1	36	4.5
pct	53%	42%	0%	3%	0%	3%	100%	-
titan	7	11	9	0	4	2	33	3.5
pct	21%	33%	27%	0%	12%	6%	100%	-

Table 45: [Hardware\_JobWait]: "Batch job wait time"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	5	11	3	0	0	2	21	4.1
pct	24%	52%	14%	0%	0%	10%	100%	-
njord	2	8	9	5	3	2	29	3.0
pct	7%	28%	31%	17%	10%	7%	100%	-
stallo	9	16	8	0	0	3	36	4.0
pct	25%	44%	22%	0%	0%	8%	100%	-
titan	9	11	6	2	0	5	33	4.0
pct	27%	33%	18%	6%	0%	15%	100%	-

Table 46: [Hardware\_JobQueues]: "Batch queue structure and job priorities"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	5	7	5	0	0	4	21	4.0
pct	24%	33%	24%	0%	0%	19%	100%	-
njord	3	11	7	3	1	4	29	3.5
pct	10%	38%	24%	10%	3%	14%	100%	-
stallo	4	18	8	1	0	5	36	3.8
pct	11%	50%	22%	3%	0%	14%	100%	-
titan	4	14	7	3	0	5	33	3.7
pct	12%	42%	21%	9%	0%	15%	100%	-

Table 47: [Hardware\_JobFlexibility]: "Batch system flexibility"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	5	7	4	0	0	5	21	4.1
pct	24%	33%	19%	0%	0%	24%	100%	-
njord	3	9	7	1	0	9	29	3.7
pct	10%	31%	24%	3%	0%	31%	100%	-
stallo	4	13	9	0	0	10	36	3.8
pct	11%	36%	25%	0%	0%	28%	100%	-
titan	1	15	9	1	0	7	33	3.6
pct	3%	45%	27%	3%	0%	21%	100%	-

Table 48: [Hardware\_JobInteractive]: "Ability to run interactively"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	3	5	2	1	0	10	21	3.9
pct	14%	24%	10%	5%	0%	48%	100%	-
njord	0	10	3	4	1	11	29	3.2
pct	0%	34%	10%	14%	3%	38%	100%	-
stallo	4	11	6	1	0	14	36	3.8
pct	11%	31%	17%	3%	0%	39%	100%	-
titan	0	13	5	1	0	14	33	3.6
pct	0%	39%	15%	3%	0%	42%	100%	-

Table 49: [Hardware\_DiskPermanent]: "Available storage capacity for permanent data (home directories)"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	7	1	8	0	1	4	21	3.8
pct	33%	5%	38%	0%	5%	19%	100%	-
njord	4	11	8	1	1	4	29	3.6
pct	14%	38%	28%	3%	3%	14%	100%	-
stallo	9	12	7	1	0	7	36	4.0
pct	25%	33%	19%	3%	0%	19%	100%	-
titan	6	10	7	3	0	7	33	3.7
pct	18%	30%	21%	9%	0%	21%	100%	-

Table 50: [Hardware\_DiskTemporary]: "Available storage capacity for temporary data"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	9	7	1	0	0	4	21	4.5
pct	43%	33%	5%	0%	0%	19%	100%	-
njord	5	15	4	1	0	4	29	4.0
pct	17%	52%	14%	3%	0%	14%	100%	-
stallo	9	14	5	0	1	7	36	4.0
pct	25%	39%	14%	0%	3%	19%	100%	-
titan	5	11	5	1	0	11	33	3.9
pct	15%	33%	15%	3%	0%	33%	100%	-

Table 51: [Hardware\_DiskPerformance]: "File system configuration and I/O performance"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	7	8	2	0	0	4	21	4.3
pct	33%	38%	10%	0%	0%	19%	100%	-
njord	6	13	3	0	0	7	29	4.1
pct	21%	45%	10%	0%	0%	24%	100%	-
stallo	6	8	10	2	0	10	36	3.7
pct	17%	22%	28%	6%	0%	28%	100%	-
titan	5	10	9	1	0	8	33	3.8
pct	15%	30%	27%	3%	0%	24%	100%	-

Table 52: [Hardware\_TapeStorage]: "Available storage capacity for archiving"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	6	2	6	0	0	7	21	4.0
pct	29%	10%	29%	0%	0%	33%	100%	-
njord	3	9	4	1	1	11	29	3.7
pct	10%	31%	14%	3%	3%	38%	100%	-
stallo	4	6	8	0	0	18	36	3.8
pct	11%	17%	22%	0%	0%	50%	100%	-
titan	5	4	5	1	0	18	33	3.9
pct	15%	12%	15%	3%	0%	55%	100%	-

Table 53: [Hardware\_TapeResponse]: "Performance of archiving/retrieval tools"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	2	2	8	0	0	9	21	3.5
pct	10%	10%	38%	0%	0%	43%	100%	-
njord	2	7	4	1	0	15	29	3.7
pct	7%	24%	14%	3%	0%	52%	100%	-
stallo	2	4	8	0	0	22	36	3.6
pct	6%	11%	22%	0%	0%	61%	100%	-
titan	1	2	6	0	0	24	33	3.4
pct	3%	6%	18%	0%	0%	73%	100%	-

Table 54: [Hardware\_TapeInterfaces]: "User interfaces for archiving/retrieval of data"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	2	2	8	0	0	9	21	3.5
pct	10%	10%	38%	0%	0%	43%	100%	-
njord	1	5	7	1	0	15	29	3.4
pct	3%	17%	24%	3%	0%	52%	100%	-
stallo	2	3	9	0	0	22	36	3.5
pct	6%	8%	25%	0%	0%	61%	100%	-
titan	1	1	6	0	0	25	33	3.4
pct	3%	3%	18%	0%	0%	76%	100%	-

Table 55: [Hardware\_NetworkSpeed]: "External network capacity to access the system and transfer data"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	7	5	6	0	0	3	21	4.1
pct	33%	24%	29%	0%	0%	14%	100%	-
njord	6	12	3	0	0	8	29	4.1
pct	21%	41%	10%	0%	0%	28%	100%	-
stallo	8	15	6	0	0	7	36	4.1
pct	22%	42%	17%	0%	0%	19%	100%	-
titan	6	13	5	0	0	9	33	4.0
pct	18%	39%	15%	0%	0%	27%	100%	-

Table 56: [Hardware\_Satisfaction]: "Overall satisfaction"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	10	6	2	0	0	3	21	4.4
pct	48%	29%	10%	0%	0%	14%	100%	-
njord	4	19	1	2	0	3	29	4.0
pct	14%	66%	3%	7%	0%	10%	100%	-
stallo	11	23	0	1	0	1	36	4.3
pct	31%	64%	0%	3%	0%	3%	100%	-
titan	7	16	6	2	1	1	33	3.8
pct	21%	48%	18%	6%	3%	3%	100%	-

## B APPENDIX for Section 6.1

Table 57: [Software\_Environment]: "User environment (shells, etc.)"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	8	8	1	0	0	4	21	4.4
pct	38%	38%	5%	0%	0%	19%	100%	-
njord	6	17	3	0	0	3	29	4.1
pct	21%	59%	10%	0%	0%	10%	100%	-
stallo	10	16	4	0	0	6	36	4.2
pct	28%	44%	11%	0%	0%	17%	100%	-
titan	10	13	6	1	0	3	33	4.1
pct	30%	39%	18%	3%	0%	9%	100%	-

Table 58: [Software\_Compilers]: "Compilers"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	7	7	0	1	0	6	21	4.3
pct	33%	33%	0%	5%	0%	29%	100%	-
njord	8	13	2	0	0	6	29	4.3
pct	28%	45%	7%	0%	0%	21%	100%	-
stallo	6	13	6	0	0	11	36	4.0
pct	17%	36%	17%	0%	0%	31%	100%	-
titan	6	12	3	1	0	11	33	4.0
pct	18%	36%	9%	3%	0%	33%	100%	-

Table 59: [Software\_Applications]: "Application software"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	4	4	1	0	0	12	21	4.3
pct	19%	19%	5%	0%	0%	57%	100%	-
njord	2	12	5	0	0	10	29	3.8
pct	7%	41%	17%	0%	0%	34%	100%	-
stallo	3	11	6	1	0	15	36	3.8
pct	8%	31%	17%	3%	0%	42%	100%	-
titan	2	9	6	1	0	15	33	3.7
pct	6%	27%	18%	3%	0%	45%	100%	-

Table 60: [Software\_Libraries]: "Programming tools and libraries"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	4	6	2	0	0	9	21	4.2
pct	19%	29%	10%	0%	0%	43%	100%	-
njord	2	11	5	0	0	11	29	3.8
pct	7%	38%	17%	0%	0%	38%	100%	-
stallo	2	11	6	0	0	17	36	3.8
pct	6%	31%	17%	0%	0%	47%	100%	-
titan	4	9	3	1	0	16	33	3.9
pct	12%	27%	9%	3%	0%	48%	100%	-

Table 61: [Software\_Perftools]: "Performance and debugging tools"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	4	3	1	1	0	12	21	4.1
pct	19%	14%	5%	5%	0%	57%	100%	-
njord	2	8	3	1	0	15	29	3.8
pct	7%	28%	10%	3%	0%	52%	100%	-
stallo	2	7	8	1	0	18	36	3.6
pct	6%	19%	22%	3%	0%	50%	100%	-
titan	0	4	5	3	0	21	33	3.1
pct	0%	12%	15%	9%	0%	64%	100%	-

Table 62: [Software\_Tools]: "General tools and utilities"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	5	6	1	0	0	9	21	4.3
pct	24%	29%	5%	0%	0%	43%	100%	-
njord	1	13	5	0	0	10	29	3.8
pct	3%	45%	17%	0%	0%	34%	100%	-
stallo	3	15	4	0	0	14	36	4.0
pct	8%	42%	11%	0%	0%	39%	100%	-
titan	2	9	6	0	0	16	33	3.8
pct	6%	27%	18%	0%	0%	48%	100%	-

## C APPENDIX for Section 6.2

Table 63: [Application\_Performance]: "Overall application performance"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	10	3	1	0	0	7	21	4.6
pct	48%	14%	5%	0%	0%	33%	100%	-
njord	5	15	3	0	0	6	29	4.1
pct	17%	52%	10%	0%	0%	21%	100%	-
stallo	7	20	2	0	0	7	36	4.2
pct	19%	56%	6%	0%	0%	19%	100%	-
titan	4	18	2	2	0	7	33	3.9
pct	12%	55%	6%	6%	0%	21%	100%	-

Table 64: [Application\_SinglePerformance]: "Sequential application performance"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	3	1	1	0	0	16	21	4.4
pct	14%	5%	5%	0%	0%	76%	100%	-
njord	2	7	2	1	0	17	29	3.8
pct	7%	24%	7%	3%	0%	59%	100%	-
stallo	5	6	4	0	0	21	36	4.1
pct	14%	17%	11%	0%	0%	58%	100%	-
titan	2	11	4	0	0	16	33	3.9
pct	6%	33%	12%	0%	0%	48%	100%	-

Table 65: [Application\_ParallelPerformance]: "Parallel application performance (scalability)"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	9	3	1	0	0	8	21	4.6
pct	43%	14%	5%	0%	0%	38%	100%	-
njord	4	16	0	1	0	8	29	4.1
pct	14%	55%	0%	3%	0%	28%	100%	-
stallo	10	14	4	0	0	8	36	4.2
pct	28%	39%	11%	0%	0%	22%	100%	-
titan	3	13	2	3	0	12	33	3.8
pct	9%	39%	6%	9%	0%	36%	100%	-

## D APPENDIX for Section 7.1

Table 66: [Support\_TechAdvice]: "Quality of support given by the technical staff"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	7	7	2	0	0	5	21	4.3
pct	33%	33%	10%	0%	0%	24%	100%	-
njord	9	15	0	0	0	5	29	4.4
pct	31%	52%	0%	0%	0%	17%	100%	-
stallo	13	12	3	0	0	8	36	4.4
pct	36%	33%	8%	0%	0%	22%	100%	-
titan	9	12	6	1	0	5	33	4.0
pct	27%	36%	18%	3%	0%	15%	100%	-

Table 67: [Support\_ResponseTime]: "Response time for queries and problems"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	7	6	3	0	0	5	21	4.2
pct	33%	29%	14%	0%	0%	24%	100%	-
njord	11	14	0	0	0	4	29	4.4
pct	38%	48%	0%	0%	0%	14%	100%	-
stallo	13	9	7	0	0	7	36	4.2
pct	36%	25%	19%	0%	0%	19%	100%	-
titan	3	12	12	1	0	5	33	3.6
pct	9%	36%	36%	3%	0%	15%	100%	-

Table 68: [Support\_ResponseUrgent]: "Response time for urgent requests (e.g. file restores, quota changes, etc.)"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	8	3	2	0	0	8	21	4.5
pct	38%	14%	10%	0%	0%	38%	100%	-
njord	9	9	2	0	0	9	29	4.3
pct	31%	31%	7%	0%	0%	31%	100%	-
stallo	11	4	4	0	0	17	36	4.4
pct	31%	11%	11%	0%	0%	47%	100%	-
titan	9	8	3	0	0	13	33	4.3
pct	27%	24%	9%	0%	0%	39%	100%	-

Table 69: [Support\_Followup]: "Follow-up to queries and reported problems"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	5	4	2	0	0	10	21	4.3
pct	24%	19%	10%	0%	0%	48%	100%	-
njord	4	13	3	0	0	9	29	4.0
pct	14%	45%	10%	0%	0%	31%	100%	-
stallo	7	11	7	1	0	10	36	3.9
pct	19%	31%	19%	3%	0%	28%	100%	-
titan	5	14	4	1	1	8	33	3.8
pct	15%	42%	12%	3%	3%	24%	100%	-

Table 70: [Support\_TimeToSolve]: "Overall time to resolve problems"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	5	7	2	0	0	7	21	4.2
pct	24%	33%	10%	0%	0%	33%	100%	-
njord	6	15	2	0	0	6	29	4.2
pct	21%	52%	7%	0%	0%	21%	100%	-
stallo	6	13	6	0	0	11	36	4.0
pct	17%	36%	17%	0%	0%	31%	100%	-
titan	2	16	7	1	0	7	33	3.7
pct	6%	48%	21%	3%	0%	21%	100%	-

Table 71: [Support\_Accounts]: "Management of project and user accounts (creation, passwords, etc.)"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	8	7	1	0	0	5	21	4.4
pct	38%	33%	5%	0%	0%	24%	100%	-
njord	10	10	2	1	0	6	29	4.3
pct	34%	34%	7%	3%	0%	21%	100%	-
stallo	10	9	2	2	1	12	36	4.0
pct	28%	25%	6%	6%	3%	33%	100%	-
titan	10	7	7	2	0	7	33	4.0
pct	30%	21%	21%	6%	0%	21%	100%	-

Table 72: [Support\_Pages]: "Quality of support pages and other end-user documentation"

resource	VERYSAT	SAT	NEUT	DIS	VERYDIS	NA	TOTAL	AVE
hexagon	3	8	5	0	0	5	21	3.9
pct	14%	38%	24%	0%	0%	24%	100%	-
njord	4	8	8	1	1	7	29	3.6
pct	14%	28%	28%	3%	3%	24%	100%	-
stallo	3	9	12	0	1	11	36	3.5
pct	8%	25%	33%	0%	3%	31%	100%	-
titan	0	13	10	3	1	6	33	3.3
pct	0%	39%	30%	9%	3%	18%	100%	-

## **E APPENDIX: Questionnaire 2009**

# Notur User Survey 2009

## Welcome

Welcome to the Notur User Survey 2009. This survey is an opportunity for the users of the Notur facilities to communicate their satisfaction with the provided resources and services, communicate experiences and opinions, and detail current and future needs.

It is important for the Notur project that the national infrastructure for computational science is developed in accordance with the needs of the users. For the Notur project, this survey is an opportunity to determine overall user satisfaction with Notur's operation, identify areas that need increased attention, and determine changing needs of the users. The results of the survey will be used in the further planning of the Notur project.

The 2009 survey is similar to the survey from last year and consists of seven parts (pages):

- User and Application information
- Notur Hardware resources
- Notur Software resources and Application performance
- Notur User support services and Resource allocation
- Overall satisfaction with Notur
- Summary
- Future needs

We invite all user groups that obtained a CPU-hour allocation on the Notur facilities in 2009 to complete the survey. The managers of the projects are encouraged to complete the survey, but may assign this task to a user that actively uses the Notur resources and has a *good overview* of the experiences and opinions of the other users in the project. We assume that the answers that we receive have been approved by the project managers and represent as much as possible the experiences and opinions of the key users in the project. We expect one completed survey per project.

Please complete the survey no later than **December 21, 2009**.

Completing the survey will take some time (15-20 minutes). We believe however that your input is important for planning future investments in resources and services. We encourage you to be as accurate as possible in your answers and provide any additional information that you think can be of use to us.

All answers will be treated confidentially. We will not reveal your identity in our publication of the final results.

If you experience any problems while filling out the survey, please report this to [sigma@uninett.no](mailto:sigma@uninett.no). Some of the forms may be wider than the default width that you use for your browser, so please adjust the width or use the scroll bar to see the complete form.

Note: *Avoid* using your web browser's built-in "Back" button as this might result in having to retype your answers in some browsers. Instead use the "Previous Page" button on the page.

1. Notur project account (format: nn\*\*\*\*k):

2. Email address:

Next page >

# Notur User Survey 2009

## User Information

1. Your name:
2. \*Which Notur hardware resources did your project use in the last 12 months?  
For the remainder of the survey, it is important that you select all the correct resources here.
  - hexagon
  - njord
  - stallo
  - titan
3. Which other Notur resources did your project use in the last 12 months?
  - User support services
  - Notur website (www.notur.no)Other:
4. How important is it for your research to have access to HPC resources?
  - Absolutely vital
  - Very important
  - Important
  - Useful
  - Not important
5. How long has your project used Notur resources?
  - Six months or less
  - 6 months–3 years
  - More than 3 years
  - Do not know
6. What is the frequency of your use of Notur resources in your project?
  - Daily
  - Weekly
  - Monthly
  - Less than once a month
  - Do not know
7. Field of science:
  - Biosciences
  - Chemistry
  - Computational Fluid Dynamics
  - Economics
  - Geosciences
  - Marine Technology
  - Material Science
  - Mathematics and Informatics
  - Medical Science
  - Petroleum Technology
  - Physics
  - Social Science
  - Structural Engineering
  - Other:

[< Previous page](#)[Next page >](#)

# Notur User Survey 2009

## Application Information

1. Please characterize your application(s).

	Yes	No	Do not know
Time (CPU-hour) demanding	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Memory demanding	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Storage demanding	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Proprietary / commercial software	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Public domain software	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Self-developed software	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Sequential application	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Embarassingly Parallel	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Shared-memory application	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Distributed-memory application	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>

2. What is important for your application(s).

	Yes	No	Do not know
High machine availability and reliability	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
High overall processing speed	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Large overall memory size	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Processor interconnect: high bandwidth	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
High system (data) security	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Large temporary storage capacity at run-time	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Large disk storage capacity for input and output data	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Large tape storage capacity for archiving	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
High external network capacity to access the system	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Availability of in-depth user support from Notur	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Availability of programming tools and libraries	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>
Availability of performance tools and debuggers	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>

3. Has your group recently performed measurements to determine how efficient your applications run on the Notur facilities, relatively to the peak performance of the machine/processors?

- Yes  
 No  
 Do not know

4. Has your group recently optimized (fine-tuned) the performance of your applications on the Notur facilities?

- Yes  
 No  
 Do not know

5. Are your most important applications suitable for 'capacity' platforms? Capacity platforms provide high job throughput and can for example be clusters of compute nodes (each with 2-4 processors, 2-8 GB memory, 80-200 GB disk) that are connected with a fast interconnect (Myrinet, InfiniBand).

- Yes  
 No  
 Do not know

6. Do you have any application that can only run (efficiently) on 'capability' architectures? A capability machine is for example a shared-memory machine with a large number of processors that share a large pool of memory and where the processors are tightly connected through a high-performance (proprietary) interconnect.

- Yes  
 No  
 Do not know

7. If you have answered yes to question 6, state the reasons why your application(s) require such platform.

7. If you have answered yes to question 6, state the reasons why your application(s) require such platform.

8. Answer only if you have answered yes to question 6. Is it possible to change the software with reasonable effort to make the application(s) run efficiently on capacity (cluster) machines? Please provide as much as possible detail.

< Previous page

Next page >

# Notur User Survey 2009

## Notur Hardware Resources

1. Please report the overall satisfaction with the Notur hardware resources used by the people in your project.

	hexagon
Machine availability (uptime)	No answer ▼
Batch job wait time	No answer ▼
Batch queue structure and job priorities	No answer ▼
Batch system flexibility	No answer ▼
Ability to run interactively	No answer ▼
Available storage capacity for permanent data (home directories)	No answer ▼
Available storage for temporary data	No answer ▼
File system configuration and I/O performance	No answer ▼
Available storage capacity for archiving	No answer ▼
Performance of archiving/retrieval tools	No answer ▼
User interfaces for archiving/retrieval of data	No answer ▼
External network capacity to access the system and transfer data	No answer ▼
Overall satisfaction	No answer ▼

## Additional Comments

1. Please give additional comments on the various resources (possibly including those that you do not currently use). For instance, you may want to comment on how well the resources match (or do not match) your needs, and why. Please specify the resource in case of any machine-specific remarks. Specify improvements that you believe are necessary as well as details about any 'Very Dissatisfied' answers given above. Future needs are addressed on page 9.

Comments on hexagon:

< Previous page

Next page >

# Notur User Survey 2009

## Notur Software Resources

1. Please report the overall satisfaction with the Notur software resources used by the people in your project.

	hexagon
User environment (shells, etc.)	<input type="text" value="No answer"/>
Compilers	<input type="text" value="No answer"/>
Application software	<input type="text" value="No answer"/>
Programming tools and libraries	<input type="text" value="No answer"/>
Performance and debugging tools	<input type="text" value="No answer"/>
General tools and utilities	<input type="text" value="No answer"/>

## Application Performance

1. Please report the overall satisfaction with the performance of the applications used by the people in your project.

	hexagon
Overall application performance	<input type="text" value="No answer"/>
Sequential application performance	<input type="text" value="No answer"/>
Parallel application performance (scalability)	<input type="text" value="No answer"/>
Max. number of processors used per job in your project	<input type="text" value="0"/>
Max. number of processors your application can use per job (in ideal circumstances)	<input type="text" value="0"/>

2. What would be the ideal resource for your project in 2009-2010? (Compose the resource that you can realistically use.)

Number of CPUs:	<input type="text" value="0"/>
Memory configuration:	<input type="text" value="Distributed"/>
Memory (GB/processor):	<input type="text" value="0"/>
Storage (TB):	<input type="text" value="0"/>
Processor interconnect:	<input type="text" value="Default (e.g., Gigabit)"/>

## Additional Comments

1. Specify improvements that you believe are necessary as well as details about any 'Very Dissatisfied' answers given above. Please specify the resource or site in case of any machine-specific remarks. Also specify software that you believe Notur should offer but currently is missing. Future needs are addressed on page 9.

< Previous page

Next page >

# Notur User Survey 2009

## Notur User Support Services

1. These questions concern the user support provided on the Notur resources. This includes helpdesk (regular) support and advanced user support. Please report the overall satisfaction from the users in your project.

	hexagon
Quality of support given by the technical staff	Noanswer ▾
Response time for queries and problems	Noanswer ▾
Response time for urgent requests (e.g. file restores, quota changes, etc.)	Noanswer ▾
Follow-up to queries and reported problems	Noanswer ▾
Overall time to resolve problems	Noanswer ▾
Management of project and user accounts (creation, passwords, etc.)	Noanswer ▾
Quality of support pages and other end-user documentation	Noanswer ▾

## Advanced User Support

1. Are you familiar with the advanced user support services provided by Notur?

- Yes, I have used it  
 Yes, but I have not used it  
 No

2. In case you have used it, were you satisfied with the results?

- Very Satisfied  
 Satisfied  
 Neutral  
 Dissatisfied  
 Very Dissatisfied  
 No answer

3. In 2003-2009, advanced user support has been project-based with project duration varying from one to ca. eight months. Based on the feedback from 2007, the advanced support activity will in 2010 also offer the possibility for longer-term (multi-year) support for specific communities. Please specify other services that advanced user support should include that will suit your needs in the best possible way.

## Resource Allocation

1. These questions concern the allocation of resources (CPU-hours) on the Notur facilities. This includes the amount of computer time granted, the distribution of the users over the compute facilities, etc. Please report the overall satisfaction from the users in your project.

	How satisfied are you?
Application procedures and feedback on applications	Noanswer ▾
Frequency and duration of allocation periods (usually 6 months)	Noanswer ▾
Available information and instructions	Noanswer ▾
Amount of resources granted at the start of a new allocation period	Noanswer ▾
Amount of extra resources granted (during allocation periods)	Noanswer ▾
Response time for applications for extra resources	Noanswer ▾

## Additional Comments

1. Specify improvements that you believe are necessary as well as details about any 'Very Dissatisfied' answers given above. Please specify the site in case of any machine- or site-specific remarks. Future needs are addressed on page 9.

< Previous page

Next page >

# Notur User Survey 2009

## Overall Satisfaction with Notur

1. Please report your **overall satisfaction** with several of the key components of the Notur project.

Available computing hardware	No answer
Available storage facilities	No answer
Available software	No answer
User support organization	No answer
User support services	No answer
Available information and documentation	No answer
Notur website ( <a href="http://www.notur.no">www.notur.no</a> )	No answer
Technical support pages ( <a href="http://docs.notur.no">docs.notur.no</a> )	No answer
Resource (CPU-hour) allocations	No answer
Overall satisfaction with Notur	No answer

2. Please report how **important** the same key components of the Notur project is to your project.

Available computing hardware	No answer
Available storage facilities	No answer
Available software	No answer
User support organization	No answer
User support services	No answer
Available information and documentation	No answer
Notur website ( <a href="http://www.notur.no">www.notur.no</a> )	No answer
Technical support pages ( <a href="http://docs.notur.no">docs.notur.no</a> )	No answer
Resource (CPU-hour) allocations	No answer
Overall satisfaction with Notur	No answer

## Additional Comments

1. Specify improvements that you believe are necessary as well as details about any 'Very Dissatisfied' answers given above. Please specify the site in case of any machine- or site-specific remarks. Future needs are addressed on page 9.

< Previous page

Next page >

# Notur User Survey 2009

## Summary

Please summarize your main experiences and satisfaction with the Notur project and other similar projects.

1. What does Notur do well? What aspects of Notur are you most pleased with?

2. What should Notur do differently? How can Notur improve? Are there specific issues where Notur can do better (services, support, etc.)?

3. How does Notur compare to other centers that you have experience with? What do other centers provide that Notur should provide as well? Please list the center(s) that you use in the comparison.

< Previous page

Next page >

# Notur User Survey 2009

## Future Needs

Please specify the needs within your project for resources in the coming 24 months. You need not answer these questions in detail if your project will cease to exist in 2009 or will have similar needs as today (but please state this below).

1. Please give a brief description of any known changes in your research activity in the coming 24 months that will affect your needs for computational and storage resources and user support. This could for example be the termination of projects, start of new projects, specific deadlines, etc.

2. Will the need for computational resources change significantly in your project in the next 24 months? If yes, please state the reason(s) and, if possible, detail the changes in your needs.

3. Will the need for storage change significantly in your project in the next 24 months? If yes, please state the reason(s) and, if possible, detail the changes in your needs.)

4. Will your project need access to hardware resources in 2009-2011 that are currently not provided by Notur? E.g., other computing architecture, other type of storage, etc. Please detail your needs.

5. Will your project have a need for services in 2009-2011 that are currently not provided by Notur? Here, one could think of for example other types of support, other types of resource allocation, services for visualization, grid-enabling of applications, etc. Please detail your needs.