

# Analysis of a survey to evaluate the performance and user satisfaction of The Norwegian High Performance Computing Consortium



**SAMBA/11/04**

Sonia Mazzi

6th May 2004



**Title:** Analysis of a survey to evaluate the performance and user satisfaction of The Norwegian High Performance Computing Consortium

**Date:** 6th May 2004  
**Year:** 2004  
**Note no.:** SAMBA/11/04

**Author:** Sonia Mazzi <mazzi@nr.no>

**Abstract:** A summary of the data collected in a recent survey to be used to evaluate the performance and user satisfaction of NOTUR, a Norwegian super-computer consortium, is presented. It comprises 113 responses out of 326 requests. In general, the levels of satisfaction are moderate to high. 84% of the respondents are mostly or very satisfied overall with NOTUR services. Users find that NOTUR compares rather well to other similar centres in the rest of the world. There is a good level of satisfaction with the user support service and its staff. The areas where users show some dissatisfaction concern waiting times for jobs to be processed and storage system and capacity. Almost half of the respondents did not rate the metacenter organisation perhaps due to lack of information or knowledge about this aspect. It would have been valuable to make an evaluation of the regional centres of NOTUR. This was not possible with the information provided. NOTUR is thought to be a very important resource for research as it links large computational resources otherwise only available to local users. Investment in a global system like NOTUR is more efficient and beneficial to the scientific community than the distribution of money between individual researchers to buy or rent their own computational resources. Work commissioned by The Norwegian Research Council.

**Keywords:** NOTUR

**Target group:** All employees

**Availability:** Open

**Project:** Diverse SAMBA

**Project no.:** 10000

**Research field:** Applied statistics

**No. of pages:** 32



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>User profile</b>	<b>2</b>
<b>3</b>	<b>Overall satisfaction</b>	<b>6</b>
<b>4</b>	<b>Specific hardware resource satisfaction</b>	<b>8</b>
<b>5</b>	<b>Storage systems</b>	<b>11</b>
<b>6</b>	<b>Software resources</b>	<b>12</b>
<b>7</b>	<b>User services</b>	<b>13</b>
<b>8</b>	<b>Technology transfer projects</b>	<b>15</b>
<b>9</b>	<b>Concluding remarks</b>	<b>16</b>
<b>10</b>	<b>Appendix</b>	<b>18</b>
10.1	Comments about GRIDUR . . . . .	18
10.2	Comments about EN/TO/TRE . . . . .	21
10.3	Comments about MAGNUM . . . . .	21
10.4	Comments about NANA . . . . .	22
10.5	Comments about SNOWSTORM . . . . .	22
10.6	Comments about FIRE . . . . .	22
10.7	Comments on storage systems . . . . .	23
10.8	Comments on software resources . . . . .	24
10.9	Comments on user services . . . . .	25
10.10	Comments on technology transfer projects . . . . .	26
10.11	Comments on what users think NOTUR does well . . . . .	26
10.12	Comments on what users think NOTUR should do differently . . . . .	28
10.13	Comments on what users think about how NOTUR compares to other such centres . . . . .	30



## 1 Introduction

The Norwegian High Performance Computing Consortium (NOTUR) is a joint mission to operate and support high performance computer equipment, software, and associated services. Partners in this consortium are: the four Norwegian universities, (Oslo, Bergen, Trondheim and Tromsø), Ceetron (a company specialising in visualisation software), The Norwegian Meteorological Institute (DNMI), the independent research organisation SINTEF, and the state-owned oil company Statoil. The project started in year 2000 for a 4-year period with a budget of about 150 Million NOK, of which about 50% was provided by the Norwegian Research Council. In 2003, the project was extended for one more year until 31/12/2004.

The Research Council of Norway is currently planning the future organisation of high performance computing (HPC) in Norway. An evaluation of NOTUR is important as a way to provide feedback to the involved parties and as input to further planning processes. The survey analysed here was implemented as part of the evaluative exercise. An international evaluation committee will meet with NOTUR management and users in Oslo during May 10.-12. 2004. The results from this survey will be important input to the committee in their work.

The survey was available on-line(<http://bdb.itea.ntnu.no/notur/survey.php>) and had 8 parts namely,

- User profile,
- Overall satisfaction,
- Specific resource satisfaction,
- Storage systems,
- Software resources,
- User services,
- Technology transfer projects,
- Summary comments.

For each section we will summarise the answers in what follows.

In total, 113 valid answers to the survey were received from a total of 326 requests to users.”

In what follows the term “user” refers to a user of NOTUR who took part in the survey. Also, “NA” means “no answer”.

User satisfaction/importance assigned was measured using an ordinal variable with the levels “Very dissatisfied/unimportant”, “Mostly dissatisfied/unimportant”, “Somewhat dissatisfied/unimportant”, “Neutral”, “Somewhat satisfied/important”, “Mostly satisfied/important”, “Very satisfied/important”. These levels were coded as levels 1 to 7 respectively.

## 2 User profile

In this section users were asked to identify themselves by giving their names and login names. The survey has a confidential character and therefore this information was not used in any way. Users were asked to give their research area (from Geosciences, Computational fluid dynamics, Material science, Marine technology, Biosciences, Chemistry, Mathematics and informatics, Physics, Economics, other).

NOTUR users were asked to list the resources they use among the hardware resources “Gridur/Embla”, “En/to/tre”, “Magnum”, “Nana”, “Snowstorm”, “Fire”, User Support services, and the NOTUR website [www.notur.org](http://www.notur.org). Users had the possibility to add other resources to the list. Other resources in use that were cited are USIT, regatta, balder, and tk-cluster.

Users were also asked to specify if they had used the resources for less than 6 months or more than 6 months and the quota they had last year.

The responses are summarised in Tables 1, 2, 3, and Figures 1 and 2.

The users had the following research areas: Biosciences 3%, Chemistry 27%, Computational fluid dynamics 19%, Economics 1%, Geosciences 27%, Marine technology 3%, Material science 1%, Mathematics and informatics 8%, and Physics 12%.

Most of the respondents had used the facilities for more than 6 months (82%).

The resource most used is Gridur/Embla (80%) and En/to/tre comes as the second most used resource (39%) followed by Magnum (27%). 33% of the users use the support services and 31% use the web site.

Quota last year	Number of respondents	Percentage of total(113)
0-50000	59	52%
50000-100000	5	4%
100000-150000	3	3%
150000-200000	3	3%
200000-250000	0	0%
250000-300000	2	2%
300000-350000	0	0%
350000-400000	0	0%
400000-450000	0	0%
450000-500000	1	1%
unlimited	1	1%
No answer	39	34%

Table 1: Quota last year.



FIELD OF SCIENCE	less than 6 months	6 months to 3 years	no answer
Geosciences	3	25	2
Percentage from Tot.Geosc.(30)	10%	83%	7%
Computational Fluid Dynamics	3	17	2
Percentage from Comp.Fl.Dy.(22)	14%	77%	9%
Material Science	0	1	0
Percentage from Mat.Sci.(1)	0	100%	0
Marine Technology	0	3	0
Percentage from Mar.Tech.(3)	0	100%	0
Biosciences	1	1	1
Percentage from Biosc.(3)	33%	33%	34%
Chemistry	3	27	1
Percentage from Chem.(31)	10%	87%	3%
Mathematics and Informatics	1	7	1
Percentage from Math.Inf.(9)	11%	78%	11%
Physics	1	11	1
Percentage from Phys.(13)	8%	85%	8%
Economics	0	1	0
Percentage from Econ.(1 )	0	100%	0
TOTAL	12	93	8
Percentage from Total	11%	82%	7%

Table 2: Length of time users have been using NOTUR.

FIELD	GRIDUR	TRE	MAGNUM	NANA	SNOWST	FIRE	SUPP	WEB	OTHER
Geosciences	26	10	0	0	0	0	10	8	1
%from Geosc.(30)	87%	33%	0	0	0	0	33%	27%	3%
Comp. Fluid Dyn	15	9	2	0	0	2	7	7	1
%Comp.Fl.Dy.(22)	68%	41%	9%	0	0	9%	32%	32%	5%
Material Science	1	1	0	0	0	0	0	0	0
% from Mat.Sci.(1)	100%	100%	0	0	0	0	0	0	0
Marine Technology	3	1	0	0	0	0	0	0	0
%from Mar.Tech.(3)	100%	33%	0	0	0	0	0	0	0
Biosciences	2	1	1	0	0	1	1	1	1
%from Biosc.(3)	67%	33%	33%	0	0	33%	33%	33%	33%
Chemistry	25	13	23	8	5	3	12	10	1
%from Chem.(31)	81%	42%	74%	26%	16%	10%	39%	32%	3%
Math and Inform	6	5	2	2	2	2	3	5	0
%from Math.Inf.(9)	67%	56%	22%	22%	22%	22%	33%	56%	0
Physics	11	4	3	1	1	3	3	4	0
%from Phys.(13)	85%	31%	23%	8%	8%	23%	23%	31%	0
Economics	1	0	0	0	0	0	1	0	0
%from Econ.(1 )	100%	0	0	0	0	0	100%	0	0
<b>TOTAL</b>	<b>90</b>	<b>44</b>	<b>31</b>	<b>11</b>	<b>8</b>	<b>11</b>	<b>37</b>	<b>35</b>	<b>4</b>
% from Total	80%	39%	27%	10%	7%	10%	33%	31%	4%

Table 3: Hardware resources use by research area.




Figure 1: Last years quota.

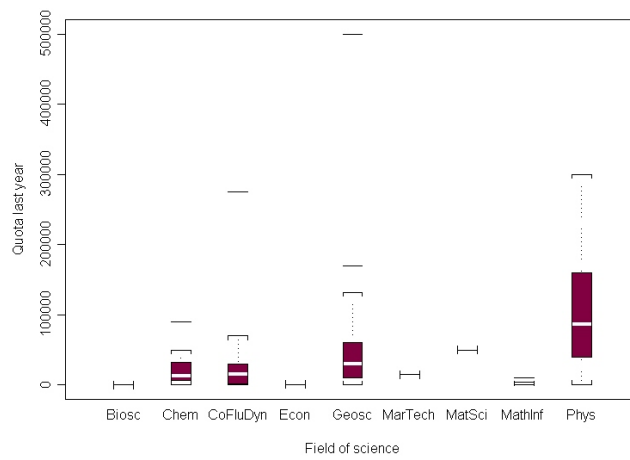


Figure 2: Boxplots of quota last year by research field.

### 3 Overall satisfaction

In this section users were asked to rate their satisfaction with NOTUR in the following 10 categories:

1. Overall satisfaction with NOTUR
2. Metacenter organisation
3. User support services
4. Notur's website
5. Allocations process
6. Available computer hardware
7. Hardware management configuration
8. Available software
9. Software maintenance and configuration
10. Mass storage facilities

Also the users's assigned importance to each of these items was asked.

The responses are summarised in Table 4. The column labelled "NA" corresponds to "no answer" and the column "AVE" is the average rating which was computed only over the answers of the respondents who actually rated the area.

The percentages are computed as percentages of the total number of respondents (113).

In general the levels of satisfaction in the aspects considered are neutral(5) to high(7). 84% of the respondents are mostly or very satisfied overall with NOTUR services.

Almost half of the respondents did not rate the metacenter organisation perhaps due to lack of information or knowledge about this aspect.

Aspect rated			very diss 1	most diss 2	some diss. 3	neutr 4	some sat 5	most sat 6	very sat 7	NA	AVE
OVERALL	Satisf	count	0	0	0	7	7	68	27	4	6.1
		%	0	0	0	6%	6%	60%	24%	4%	
SATISF	Import	count	0	0	0	0	3	16	85	9	6.8
		%	0	0	0	0	3%	14%	75%	8%	
METACEN	Satisf	count	0	0	0	24	5	21	12	51	5.3
		%	0	0	0	21%	4%	19%	11%	45%	
ORGANIS	Import	count	0	0	0	0	23	26	15	49	5.9
		%	0	0	0	0	20%	23%	13%	43%	
USER SUPPORT SERVICES	Satisf	count	0	0	1	10	6	40	47	9	6.2
		%	0	0	1%	9%	5%	35%	42%	8%	
IMPORT	Import	count	0	0	0	0	1	20	81	11	6.8
		%	0	0	0	0	1%	18%	72%	10%	
NOTUR WEBSITE	Satisf	count	0	1	6	17	14	45	11	19	5.4
		%	0	1%	5%	15%	12%	40%	10%	17%	
IMPORT	Import	count	0	0	0	0	12	60	20	21	6.1
		%	0	0	0	0	11%	53%	18%	19%	
ALLOCAT PROCESS	Satisf	count	1	0	6	15	13	27	19	32	5.4
		%	1%	0	5%	13%	12%	24%	17%	28%	
IMPORT	Import	count	0	0	0	0	6	26	44	37	6.5
		%	0	0	0	0	5%	23%	39%	33%	
AVAILABLE COMPUTER HARDWARE	Satisf	count	0	1	6	5	13	36	32	20	5.9
		%	0	1%	5%	4%	12%	32%	28%	18%	
IMPORT	Import	count	0	0	0	0	1	6	82	24	6.9
		%	0	0	0	0	1%	5%	73%	21%	
HARDWARE MANAGEM CONFIG	Satisf	count	0	0	1	7	8	37	26	34	6.0
		%	0	0	1%	6%	7%	33%	23%	30%	
IMPORT	Import	count	0	0	0	0	2	16	57	38	6.7
		%	0	0	0	0	2%	14%	50%	34%	
AVAILABLE SOFTWARE	Satisf	count	0	0	1	17	7	34	29	25	5.8
		%	0	0	1%	15%	6%	30%	26%	22%	
IMPORT	Import	count	0	0	0	0	8	31	45	29	6.4
		%	0	0	0	0	7%	27%	40%	26%	
SOFTWARE MAINT	Satisf	count	1	0	4	17	9	27	20	35	5.5
		%	1%	0	4%	15%	8%	24%	18%	31%	
IMPORT	Import	count	0	0	0	0	5	26	45	37	6.5
		%	0	0	0	0	4%	23%	40%	33%	
STORAGE FACILIT	Satisf	count	4	2	10	12	9	25	23	28	5.2
		%	4%	2%	9%	11%	8%	22%	20%	25%	
IMPORT	Import	count	0	0	0	0	7	25	48	33	6.5
		%	0	0	0	0	6%	22%	42%	29%	

Table 4: Overall satisfaction.

## 4 Specific hardware resource satisfaction

In this section users were asked to rate their satisfaction with the specific hardware resources Gridur/Embla, En/to/tre, Magnum, Nana, Snowstorm, and Fire in the following aspects:

1. Overall satisfaction
2. Uptime
3. Batch job wait time
4. Batch queue structure
5. Ability to run interactively
6. Disk configuration and I/O performance
7. Amount of available computer time
8. Amount of available storage capacity

Also, users were asked to state:

1. Maximum number of processors used per job
2. Maximum number of processors your code can effectively use per job.

Some respondents provided comments on the specific hardware resources mentioned. These comments are listed in the Appendix.

The answers are summarised in Tables 5 and 6 and in Figures 3 and 4.

The average ratings are computed only using the actual responses (i.e. not considering the NAs). The percentages are percentages of respondents who rated the aspect of the particular hardware resource.

Aspect rated		Hardware resource					
		Gridur/ Embla	En/To/Tre	Magnum	Nana	Snow Storm	Fire
Overall satisf	Ave %resp	5.7 72%	6.2 34%	6.1 22%	6.2 8%	4.9 7%	5.2 5%
Uptime	Ave %resp	5.4 65%	6.5 30%	6.1 21%	6.7 9%	5.3 5%	6.4 4%
Batch job wait time	Ave %resp	4.8 68%	4.9 34%	5.2 21%	5.5 9%	4.3 5%	5.4 4%
Batch queue structure	Ave %resp	5.2 58%	5.5 29%	5.4 19%	6.0 9%	5.0 5%	4.0 4%
Ability run interactive	Ave %resp	5.1 53%	5.0 29%	4.9 17%	5.1 8%	4.5 5%	5.2 4%
Disk config I/O perf	Ave %resp	5.4 52%	5.6 25%	5.6 14%	6.2 5%	5.3 3%	5.2 4%
Amount av comp time	Ave %resp	5.7 62%	6.0 28%	5.8 20%	6.6 6%	6.7 3%	6.2 4%
Amount av storage cap	Ave %resp	5.3 61%	5.1 28%	6.1 16%	4.8 5%	4.7 3%	4.6 4%

Table 5: Ratings of hardware resources.

Maximum number of processors used per job						
Summary statistics	Hardware resource					
	Gridur	En/to/tre	Magnum	Nana	Snowstorm	Fire
Min	1	1	1	1	1	1
Mean	39	11	6	3	3	12
Median	16	8	4	1	1	8
Max	280	32	16	16	16	32
% of NA's	36%	73%	83%	93%	95%	96%
Maximum number of processors a code can effectively use per job						
Min	1	1	1	1	1	1
Mean	154	44	28	46	57	16
Median	48	16	6	2	6	8
Max	1024	400	256	256	256	64
% of NA's	50%	76%	86%	95%	96%	95%

Table 6: Maximum number of processors used and needed.

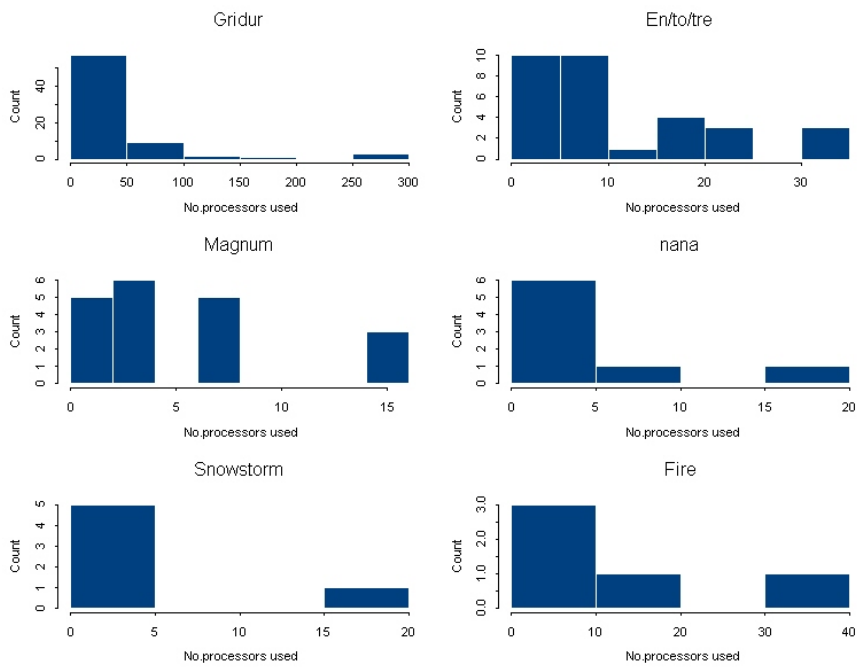


Figure 3: Maximum number of processors used per job by machine.

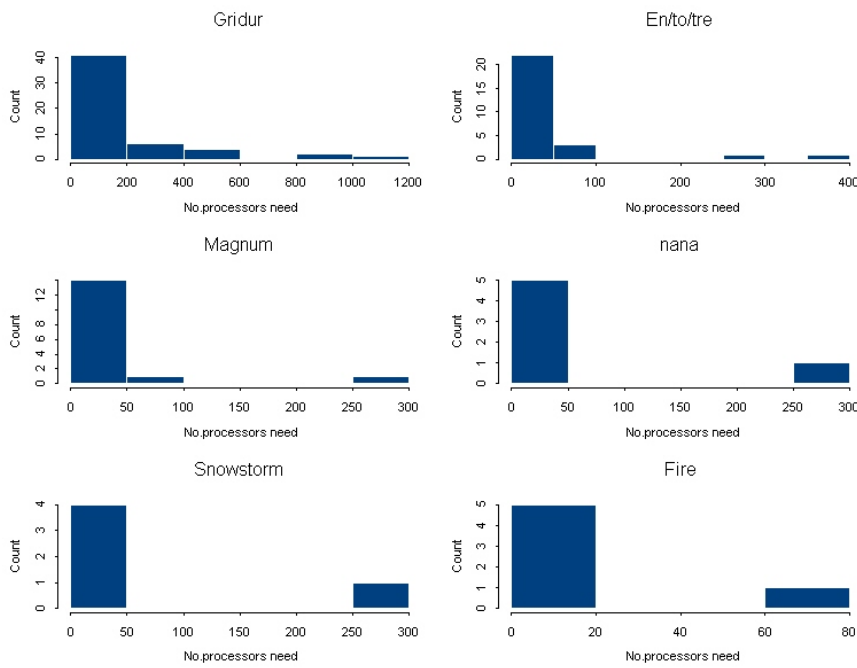


Figure 4: Maximum number of processors the code can effectively use per job by machine.



## 5 Storage systems

Users were asked to rate their satisfaction on the following aspects of storage systems:

1. Overall satisfaction
2. Uptime
3. Reliability
4. Performance
5. System response time
6. User interface (scp, sftp, etc.)

The answers are summarised in Table 7.

Note that the percentages of non-respondents is quite high (above 35%) perhaps due to the rather technical level of the question.

Respondents also provided comments on storage systems which are listed in the Appendix.

The average ratings are computed excluding NAs.

The average ratings are in all cases above “neutral” (5).

Summary stats	Aspects rated					
	Overall satisf	Uptime	Reliable	Performance	Response time	User interface
Ave	5.9	6.2	6.3	6.1	5.5	6.0
% NA's	35%	42%	39%	39%	42%	37%

Table 7: Ratings of storage systems.

## 6 Software resources

In this section users were asked to rate their satisfaction with the software resources in each of the six hardware resources Gridur/Embla, En/to/tre, Magnum, Nana, Snowstorm, and Fire in the following aspects:

1. User environment
2. Fortran compilers
3. C/C++ compilers
4. Application software
5. Programming libraries
6. Performance and debugging tools
7. General tools and utilities

The answers are summarised in Table 8.

Averages are computed excluding NAs and percentages are percentages of respondents who rated the aspect in the specific hardware resource.

Also, users were able to provide comments about software resources, suggested improvements, future needs, etc. The comments are listed in the Appendix.

Aspect rated		Hardware resource					
		Gridur/ Embla	En/To/Tre	Magnum	Nana	Snow Storm	Fire
User Environm	Ave %resp	6.0 59%	6.3 27%	6.1 15%	6.0 6%	5.6 4%	5.3 3%
Fortran compilers	Ave %resp	6.0 44%	6.3 24%	5.6 9%	6.0 4%	5.0 2%	6.3 3%
C/C++ compilers	Ave %resp	5.6 23%	6.0 12%	5.0 5%	6.0 4%	5.0 4%	6.5 2%
Applic software	Ave %resp	5.6 32%	5.6 18%	6.0 11%	6.0 4%	6.7 3%	5.3 3%
program libraries	Ave %resp	5.8 33%	5.8 14%	5.8 7%	6.0 3%	5.0 4%	6.0 2%
Performance debug tools	Ave %resp	5.0 32%	5.3 14%	5.7 6%	6.0 3%	6.5 2%	7.0 1%
Gral tools utilities	Ave %resp	5.5 31%	5.6 14%	6.0 6%	5.6 4%	4.7 3%	6.0 1%

Table 8: Ratings of software resources.

## 7 User services

The questions in this section pertained to NOTUR User Services: Helpdesk, Advanced user support, and resource allocation.

Users were asked to rate their satisfaction on the following aspects:

1. Quality of technical advice from consultants
2. Timely initial response to consulting questions
3. Follow-up to initial consulting questions
4. Amount of time to resolve your issue
5. If you have received advanced user support: quality of technical advice
6. Response to special requests (e.g. file restores, quota changes, etc.)
7. Account Support Services (account creation, passwords, etc.)
8. Amount of computer time allocated, regular allocation periods
9. Amount of extra computer time allocated (between regular periods)
10. Extra computer time, response time on applications

Answers are summarised in Table 9.

Average ratings were computed excluding NA's.

All average ratings are above "mostly satisfied" (6).

Respondent provided comments on NOTUR User Services. These are listed in the Appendix.

Aspect rated		very diss 1	most diss 2	some diss. 3	neutr 4	some sat 5	most sat 6	very sat 7	NA	AVE
QUALITY ADVICE	count	0	0	2	3	8	25	46	29	6.3
	%	0	0	2%	3%	7%	22%	41%	26%	
TIMELY RESPONSE	count	0	0	3	6	9	20	49	26	6.2
	%	0	0	3%	5%	8%	18%	43%	23%	
FOLLOWUP	count	0	0	2	6	6	23	44	32	6.2
	%	0	0	2%	5%	5%	20%	39%	28%	
TIME TO SOLVE	count	0	1	4	4	10	25	37	32	6.0
	%	0	1	4%	4%	9%	22%	33%	28%	
QUAL.ADVISE ADVANCED	count	2	0	2	2	3	10	23	71	6.0
	%	2%	0	2%	2%	3%	9%	20%	63%	
SPECIAL REQUEST	count	0	0	1	2	3	19	28	60	6.3
	%	0	0	1%	2%	3%	17%	25%	53%	
ACCOUNT SERVICE	count	0	0	5	4	3	20	33	48	6.1
	%	0	0	4%	4%	3%	18%	29%	42%	
TIME ALLOCATED	count	0	1	2	3	5	30	22	50	6.0
	%	0	1%	2%	3%	4%	27%	19%	44%	
AMT.EXTRA COMP.TIME	count	0	0	2	6	3	19	27	56	6.1
	%	0	0	2%	5%	3%	17%	24%	50%	
APPLY EXTRA COMP.TIME	count	0	0	0	7	6	13	23	64	6.1
	%	0	0	0	6%	5%	12%	20%	57%	

Table 9: user services ratings.

## 8 Technology transfer projects

This section referred NOTUR Technology Transfer Projects (TTPs). Perhaps few respondents participated in these TTPs and hence the high number of NAs.

Users were asked to rate their satisfaction in the following aspects of TTPs:

1. Overall satisfaction
2. Cooperation with other personnel
3. Cooperation with other institutions
4. project achievements
5. Project funding
6. Available personnel time

Answers are summarised in Table 10.

Comments provided by users in this section are listed in the Appendix.

Aspect rated		very diss 1	most diss 2	some diss. 3	neutr 4	some sat 5	most sat 6	very sat 7	NA	AVE
OVERALL	count	0	0	0	1	3	10	5	94	6.0
SATISF	%	0	0	0	1%	3%	9%	4%	83%	
COOPER	count	0	0	0	1	3	3	6	100	6.1
PERSON	%	0	0	0	1%	3%	3%	5%	88%	
COOPER	count	0	0	0	1	3	3	6	100	6.1
INSTIT	%	0	0	0	1%	3%	3%	5%	88%	
PROJECT	count	0	0	0	1	2	4	7	99	6.2
ACHIEVM	%	0	0	0	1%	2%	4%	6%	88%	
PROJECT	count	0	0	2	2	1	3	2	103	5.1
FUNDING	%	0	0	2%	2%	1%	3%	2%	91%	
PERSON	count	0	0	0	3	1	3	2	104	5.4
TIME	%	0	0	0	3%	1%	3%	2%	92%	

Table 10: Technology transfer projects ratings.

## 9 Concluding remarks

This report presents a summary of the responses from the survey recently carried out to evaluate performance and user satisfaction of NOTUR. The author will not attempt to draw strong conclusions to evaluate NOTUR from these summaries. This is left to the commissioners of this work.

However, it has been observed that, in general, **there seems to be a moderate to high level of satisfaction with the service.**

Although in some instances the level of non-respondents is high, I can imagine that those responses that were obtained, although few, are very valuable as they are probably coming from users with high expertise in super-computing.

Some respondents have provided extensive and apparently insightful comments on different aspects of NOTUR (listed in the Appendix), which the actual providers of the service should read attentively. I think that these comments will provide invaluable feedback to the system administrators.

In general the **comments show a good level of satisfaction with the user support service, availability, responsiveness and helpfulness of the user support service staff.** Several users praise the non-bureaucratic character of the administration of the service.

**Users find that NOTUR compares rather well to other similar centres in the rest of the world.** It is praised that NOTUR is more easily accessible/available to researchers than other such facilities.

However, as one user notes, there is a slight deficiency in this survey: respondents should have identified which center of NOTUR they use the most or are associated with. Some users show dissatisfaction in the same areas where others show complete satisfaction and this variation could be due to, as a user suggest, regional variations in the level of support, expertise of personnel at different centres, etc. **It would have been interesting to see if there are notable differences in the way different NOTUR centres function.**

The hardware resource mostly used is Gridur/Embla and respondents have provided several comments on the functionality and administration of this resource. Respondents report much “down-time”, and high dissatisfaction with the queue system and the fact that jobs seem to be suspended too often and for too long. Users report that it is specially in the last four months or so that the time to process jobs has increased. These **delays processing jobs contradict the very principle of super-computing and therefore should be addressed with much attention.** It is reported that upgrades in the resource cause too much unwanted instability. Perhaps these upgrades should be better planned so that they do not disrupt ongoing projects.

Respondents show dissatisfaction with the storage systems. They point to **inadequacy of the current storage system which is based solely on disk storage.** This causes that files are deleted too soon in some instances. **Several users point to the need to use a tape robot system at all NOTUR centres.** Users seem to find their storage space too limited to carry out their projects effectively. This area needs serious consideration by administrators.

To conclude, as one user points out, **NOTUR is a very important resource for research since it links large computational resources otherwise only available**

to local users. Investment in a global system like NOTUR is more efficient and beneficial to the scientific community than the distribution of money between individual researchers/groups to buy or rent their own computational resources.

## 10 Appendix

In this Appendix, comments provided by respondents are listed. These comments are presented verbatim without modifications. Irrelevant comments are not listed.

### 10.1 Comments about GRIDUR

30 comments

1. “The machine crashes a bit too often. The number of processors is wonderful, but sadly our software is poorly implemented. Only due to the huge number of processors I find it justifiable to use this machine intensively.”
2. “Problems with queue system and possibility to run large jobs (i300–400)”
3. “Storage of ocean model output and forcing fields is a real problem. There would be a huge improvement in my productivity if a tape robot system could be re-established.”
4. “Parallel speedup is difficult to check because processors are shared with other jobs.”
5. “Interactive response time has been extremely large lately. This is really annoying and time consuming. Queue situation is hopeless.”
6. “Have been much downtime, not used it for a while”
7. “This facility matches all my needs. The only problem is that I have not been given the amount of CPU time which my project requires.”
8. “Overall I am very satisfied with these computer, especially compared to other machines I have access to in the US. They are quite fast, and the queue time on average over the last few years have been quite good. However, the last month or so it has not been quite as good. If this is only a fluke, I won’t say anything, but I do hope it won’t be permanent. Also, there is one quite annoying thing — I can only log in via ssh from local computers at the University of Oslo, not from home or from JPL/CalTech, USA, or anywhere else. That means I have to first ssh into some computer in Oslo, and then into Gridur/Embla. It works, but it is somewhat annoying.”
9. “Better facilities for code debugging, and better response times for debugging in totalview”
10. “The machines go down way too often!”
11. “The single queue is very slow. Though I was grateful to be allowed to use normal queue the months before handing in my thesis.”
12. “Max. no of processors used is a typical no. Max no. of effective processors depends on data size, and could be larger given a larger machine.”



13. “Mixture of codes. Some can use 512 processors effectively but most flatten out around 32 processors. Work is in progress to increase this number for the future. No major problems except for stability in periods and long waiting times for jobs using more than 48 processors (which is why that number figures in the table). Good for codes that parallelize on many processors because of large amounts of time available.”
14. “Difficult to answer. I am account responsible for the users in the RegClim project. But in climate research computer resources are a limiting factor.”
15. “My usage goes up and down. I had a period with intense use of gridur/embla in December 2003, it worked out very well. I presently have another such period, and compared with my experience from 2003–12, the situation is dramatically worse (what used to take 1 day now takes two weeks). I’m presently looking elsewhere for available HPC facilities, but the situation is very difficult. I have a deadline on a NFR project on July 1 this year, but given the present turn-around time of my jobs, I almost certainly can’t meet that deadline.”
16. “Most runs are tests and spin up og wave- and ocean model code for operational forecast runs at met.no. Some runs are hindcast simulations in connections with external projects. For operational runs, computer efficiency and up time are essential. gridur/embla mostly match those needs.”
17. “In some periods we need large amounts of calculations to be done in a short period, and this can efficiently be done on embla/gridur. This has been very valuable for us.”
18. “They mostly work ok for me. Some times too long time to wait to get a job through the batch system.”
19. “Very bad storage capacity! Otherwise I am very satisfied.”
20. “I have not done any computations myself, the computations are done by a cooperating researcher”
21. “Unødvendig å være så streng når regnekontoen er oppbrukt. Noen ganger trenger jeg bare å teste programmet mitt om det virker på ulike plattformer, og da hadd e det vært fint å kunne kjøre små jobber selv om det ikke er mer regneid igjen.”
22. “En ting jeg ikke forstår er at når man har en maskin som fungerer bra såskal man alltid oppgradere den. Etter slike oppgraderinger har jeg inntrykk av at maskinen blir ustabil for en tid, med i/o problemer og crashing. En annen ting jeg vil sette fingeren på er køsystemet. At noen kan ha 5,6,7 jobber gående samtidig mens andre står i kø med en single jobb. Ellers syns jeg ting fungerer bra, og man får god hjelp hvis man trenger det. ”
23. “The batch system typically offers CPU-time to many jobs on average, but a significant number of them are suspended at any given time. This strategy is puzzling to me, and also makes the planning of projects and distribution of jobs between the

different NOTUR centers difficult. We basically never know when the dridur/embla jobs are going to finish. It is OK to be queued because jobs submitted earlier or with a higher priority are currently running. From the queue-list you may guess when a given job may start to run. If it seems to be taking too long, you can always delete it and submit it elsewhere. But when a job basically is in; i.e., has been offered CPU-time, but is suspended perhaps 22 hours per day, when will it finish? Will it be suspended on average 20, 22, or perhaps 24 hours per day in the weeks to come?"

24. "The machine is heavily focused on parallel jobs, which recently has not been my main focus. However, each individual processor is too slow to be of much interest, and wait-time for highly parallel jobs (to compensate for slow processors) appear too high."
25. "I don't like the queue system, all the jobs start in a very few time, but they are suspended for very long time. Maybe it is useful to improve the time distribution."
26. "Invaluable source for my work results. Best available machines for me currently, because of its relatively large memory accessibility. Although the IBM Regatta computer in Bergen is faster, it does not have as large memory available, and this is vital to my applications"
27. "I find security aspect somewhat over stressed. Life would be easier if one could open ports without rebooting (have I misunderstood?), access the disks form other machines, mount external disks etc. Very Very helpful staff. Thank you!"
28. "Too many stops. Too little info on what's going on"
29. "1) System Upgrade has NOT been satisfactory. The users should not be offered an upgraded system for use before it is at least working as properly as the preious version. After mo re than two weeks, the system worked on gridur. On embla, the system still have important problems after more than a month! 2) User support for optimization of professional program code to actual hardware configuration has not functioned according to expectations. We have been granted a user support project for adjusting the program system to the hardware configuration. After 3 months hardly anything has happened. It appears to us that there are local educational aspects involved that are not in the interest of the users, but we may be wrong. On earlier occasions when we have had assistance from expertise in Trondheim, we have been satisfied. It seems that the problems now have been delegated to local expertise at UiO that does not have the necessary competence. "
30. "The queues structure: Most of my jobs are processed with only one processor and should be placed in the single queue. Embla which supports the single queue does however not support FLUENT, so all my jobs have to be placed in the normal queue on Gridur."

## 10.2 Comments about EN/TO/TRE

9 comments

1. “It is, in my view, virtually impossible to do any serious numerically-intensive computing on this system now. There is no point running my large models with anything less than 16 processors and it is impossible to access that many processors these days. I am very pleased with the tape robot system.”
2. “I am new on the system and just beginning my computations in bio-geochemical ocean modelling. My experience with the staff so far was very good.”
3. “This machine is very good, but is so overloaded that it has not been possible to extract the granted CPU time. With the current user load it can only be used for small jobs. For example, 8-processor jobs with wall-clock limits on a few hundred hours will be waiting in queue for weeks. This cannot be associated with super-computing.”
4. “This machine suits particularly well our applications in ocean forecasting. It is further very efficiently managed and reliable. Special thanks for the advanced user support.”
5. “Nice machine. Only problem is load and therefore possibility of getting large allocations of time.”
6. “Excellent machine, that has been a great asset to our research recently. Fast processors, well-functioning both for single- and multiprocessor jobs. Somewhat long wait-time for many-processor jobs.”

97

7. “Miss more info on submitting jobs. I.e. the amount of Consumable Memory available. Some problems running large jobs, probably due to memory available.”
8. “Superb support staff! Clever people from that group can work magic with the performance of your code.”
9. “It takes a very long time to get monster jobs started, which I think is a big problem. Also I have a problem with the fact that files on work can get deleted after only 3 days. This is too soon, especially when you have jobs that are running for more than 3 days. The thing I usually do is to move result files to migrate when the job has finished, but by then your files can already be deleted.”

## 10.3 Comments about MAGNUM

4 comments

1. “Great machine, fabulous up-time. Especially the speed is amazing and I wish we had more of such fast processors...”

2. “Long queues.”
3. “Nice machine. Some initial problems after the upgrade to get the codes running and somewhat slow response from support solving the problem. Good for running codes that do not parallelize on more than a few processors because of the good performance per processor.”
4. “After the upgrade and the problems associated with this, the machine has been functioning well, with reasonable wait-times in queue. Have not tried to make a parallel code on this machine.”

## 10.4 Comments about NANA

2 comments

1. “Somewhat long queue. Used by us a lot. Should have been bigger? To bio-inorganic is it very important.”
2. “Despite it’s old age, still performing reasonably well, probably due to a somewhat lower user number than the other machines(?). Very reliable, and my favourite pick for single processor or multiprocessor jobs that need to run for a very long time (1 month).”

## 10.5 Comments about SNOWSTORM

3 comments

1. “Waiting for adf. Very good to have available.”
2. “For some reason, snowstorm refused to run (MPI) code which it compiled itself. I am sure it could have been worked out, but I was in a hurry at the time and the error messages were not very informative.”
3. “Still awkward with the different processors. Has been on-and-off with respect to availability of some of the itanium machines.”

## 10.6 Comments about FIRE

2 comments

1. “Problems configuring stuff.”
2. “The interconnect is not good enough. My application don’t scale.”

## 10.7 Comments on storage systems

13 comments

1. "Tape robot systems should be established at all NOTUR centres. The work areas on disks seem large enough, but it is difficult moving data from large simulations from the work area to more permanent storage"
2. "Could have been more space on nana?"
3. "As mentioned above, I can only log in from computers at the University of Oslo. Another thing is that the disk usage policy is quite aggressive — if one doesn't access a file on the scratch drive within two weeks, it is deleted. This has been a problem in some cases, but I do understand that it is necessary with some restrictions."
4. "My needs are taken care of by NoSerC, and their solution works is satisfactory."
5. "Would assume larger disk space was available for user data."
6. "Do not have strong requirements for storage and am therefore easily satisfied."
7. "Difficult to answer. I am account responsible for the users in the RegClim project. But in climate research computer resources are a limiting factor. Specially lack of storage capacity has been a problem."
8. "Uses only disk. Not much use of the tape Robot"
9. "The larger the disk, the better...."
10. "Is there anything else than the work-disk?"
11. "I have not done any computations myself, the computations are done by a cooperating researcher"
12. "The usual disk-quota is rather small compared to the size of the result files of a typical contemporary research project. The individual research groups thus must take pains to transfer and handle these data elsewhere prior to analyzing them and writing up publications etc.. The small limits (typically 1–2 GB per user) are surprising given the dramatic fall in harddisk prices in recent years. Some sites offer migration solutions (tape robots etc), but the considerable access times exclude such solutions during the active phase of a project when individual result files may be accessed on a regular basis. The migration systems are thus better suited as a backup solution when the projects are basically finished. "
13. "We mainly use the Tape robot system installed bay the NoSerC-project, because we believe what NOTUR can offer is inadequate. Our main problem is the limited amount of space allocated for storage."

## 10.8 Comments on software resources

17 comments

1. “We should have Turbomole, which is much more efficient than Gaussian/ADF and is even cheaper when ordered through the research group of Prof. Alrich at University of Karlsruhe. It came to my attention that Turbomole was considered to be expensive, as offers were made by Accelrys, but they do not sell the software, but a user-interface.”
2. “Matlab should be available on gridur/embla. Matlab should be updated with a newer version on en/to/tre. netCDF libraries and utilities should be installed on gridur/embla. netCDF support should be improved on on/to/tre. The NCAR graphics package should be installed on gridur/embla and en/to/tre. grib (met. data) utilities and libraries should be placed on gridur/embla and en/to/tre.”
3. “One thing I would like, is the possibility to log into Gridur/Embla via Emacs remotely. I guess what prevents me from doing so is the same feature as prevents me from ssh into these machines in the first place. Because of this I have to run Emacs in a terminal, which is not as convenient as I would wish.”
4. “The problem is that quite often following system “upgrades” the machine is very unstable, and sometimes has extremely poor performance. A good example of this happened in late March, and it took several weeks until the problem was solved (Good Friday on gridur). In the meantime, the machine was more than a factor 10 slower than normal. Why wasn’t this problem acknowledged and resolved earlier?”
5. “Inconsistencies in C++ libraries not resolved and seem to receive low priority when reported.”
6. “All codes developed by us and written in Fortran. Debugging done on local systems. Therefore no answers on the other questions. Same needs for the future: stable systems with good Fortran (and C++ for some others in the group) compilers, no special needs for other software resources.”
7. “Difficult to answer. I am account responsible for the users in the RegClim project and not an active user myself.”
8. “-ssrun difficult/impossible to use on many nodes -totalview often crashes”
9. “We need netcdf library.”
10. “I miss a good, quick and easy to use debugger. (..or maybe I do not know about it..)”
11. “I have not done any computations myself, the computations are done by a cooperating researcher”

12. “Jeg trenger mange ulike kompilatorer for å teste at koden virker med de ulike kompilatorene. Ingen problemer så langt med dette.”
13. “More support for C++ compiler generally. More support for I/O libraries such as HDF/netCDF.”
14. “The Gaussian program package is probably one of the most important applications on the NOTUR machines, in terms of consumed CPU time. Gaussian 03 offers shared-memory parallelization of most of the important routines, but the speedup seldomly warrants the use of more than, say 4 processors. The poor parallelization of the shm-version of Gaussian indicates that this program is not very well suited to the (predominantly parallel) NOTUR hardware. The strong position of Gaussian among chemists suggests that this program will nevertheless be a large consumer of CPU time in the years to come. One should rather look for possibilities for adjusting the hardware, or perhaps better: consider the possibility of obtaining a distributed memory version of Gaussian. Thus, one would have to consider purchasing a license for the Linda parallel execution environment. This is probably expensive, but might still pay off given the number of users and the profile of the NOTUR hardware.”
15. “Idl”
16. “You don’t really expect us to have qualified opinion on this? All I can say is that the compilers seems to without serious bugs,”
17. “The Fortran compiler should be better equipped with debuggers.”

## 10.9 Comments on user services

13 comments

1. “Somehow, I find it easy to miss application deadlines and I am never quite sure whether or not my application has gone through...”
2. “Very service-minded. Very good. Keep up the good work. Special thanks to the guys in Tromsø who shows up personally if necessary.”
3. “Vegard Eide is very helpful and knowledgeable.”
4. “Simplify getting accounts on additional resources within the same grant.”
5. “Name confusion between user account and quota account.”
6. “Takes too long to create new accounts (several days to a week)”
7. “Difficult to answer. I am account responsible for the users in the RegClim project and not an active user myself.”
8. “I use the account n2857k. This is a joint account for more people at the Norwegian Meteorological Institute. I therefore have not applied for quota myself.”

9. "More expertise"
10. "There are large differences in the quality of user support between the different centers. The form above thus should have distinguished between the NOTUR centers."
11. "Very satisfied in general. I am impressed with the responsiveness and the consideration shown by the employees. Also with the flexibility in allocating extra CPU time when needed at vital moments. The allocation process, with having to set a certain assumed CPU quota every half year though, I don't like."
12. "I consider myself an advanced user. Thus the only support I've had is "advanced user support" from Bergen. All my answers are related to this experience"
13. "Please see comment No. 2 above. We were allocated a special project on optimization of professional computer code to hardware configuration. The progress of this project is dissatisfactory."

### 10.10 Comments on technology transfer projects

3 comments

1. "-started and finished too late"
2. "We are in the initial stage of the project, so we have not obtained any results yet."
3. "The political constrains (No institution larger then 50%, "Vertical integration") was a strait jacket."

### 10.11 Comments on what users think NOTUR does well

33 comments

1. "It appears at first sight to be very rigid and inflexible with its webforms, application deadlines and somewhat unclear questions in the very same forms. Despite of this, Notur is very generous, helpful and grants the benefit of doubt."
2. "Information about coming events is OK, but information about ongoing problems might be better."
3. "Easy and fast access to computer time Good service"
4. "The support is great!"
5. "uptime, stability, etc."
6. "Metacenter Computer time allocation Operations"
7. "Notur has done an excellent job of making advanced super-computing facilities available to the Norwegian academic and research community."



8. "I am new on the system. The staff is friendly and helpful. Maybe in 6 months I can give you a more detailed response."
9. "Notur run the facilities very well. The availability of support is very good."
10. "In my opinion the user support offered by the technical staff is excellent."
11. "I can only compare to other machines I have access to, and the main advantages of Notur for me, is short queue time (although not as good lately as it has been) and "unlimited" memory. In fact, I prefer to use these machines over US-based ones whenever possible, just for these reasons."
12. "I am very happy about the fact that login shells are more than responsive enough to allow interactive sessions for debugging."
13. "1) Flexibility in connection with need for extra computer time, higher priority, etc.  
2) User support in Trondheim (in some cases)."
14. "Synergy: Large computational resources that it would be very inefficient to have locally run and located closer to the users. It would be meaningless to distribute the money used in NOTUR to the researchers themselves in order for them to buy/rent their own resources. Very good cooperation among the people running the four sites on a daily basis. Hardware of good quality. "
15. "Support at all places."
16. "I'm most pleased with the user support. Very helpful."
17. "User support er effektivt og bra."
18. "Hardware resources, general operations, uptime, time allocation"
19. "Allocation of HP computer resources are of crucial importance for climate modelling."
20. "User support"
21. "Computational resources. Up time."
22. "Provide stable and efficient computer resources and expertise"
23. "User support The express queue works fine for small jobs and debugging."
24. "Run hardware. Maintain software. The idea of a national center leads to better exploitation of the hardware and a necessary concentration of competence."
25. "Very seldom down time – stability of the systems good information and time before down time"
26. "Adequate hardware and software most of the time. Application for CPU time is not overly time-consuming and difficult."

27. "Notur has provided stable, reliable access to high-performance computing resources. Different architectures are available, providing us with the choice of choosing the architecture most suited for any given application. The metacenter structure, with distributed computer service centers, allow for local expertise at each University site to help with specific problems, often of a highly technical kind. The non-bureaucratic application procedures which gives you a two application rounds access to computer resources at several different sites."
28. "They can give advices as soon as possible."
29. "Brukerstøtten i forhold til åfåpåplass den programvaren jeg trenger har vært bra. Kunne tenke meg litt mere mulighet til åkomme i kontakt med andre brukere av programmer for åfåråd og tips om hvordan ågjøre ting."
30. "Short response time on administrative tasks."
31. "The service of the employees."
32. "Helpful initial user support. Normally, the system is user-friendly. Some central staff member are very competent."
33. "General availability of large computer resources covering a range of different platforms/architectures. Stability and reliability."

## 10.12 Comments on what users think NOTUR should do differently

28 comments

1. "Please send out email reminders of application deadlines and email confirmation successfully submitted quota applications."
2. "The system with applications for CPU quotas each half year should be replaced by applications without deadlines, not necessarily for half year periods. Batch queues are often very long, which can be a bit discouraging. Could it be possible to get realistic estimates of when a job submitted to a batch queue would start? Or maybe the current hardware resources simply are too small for the CPU demands."
3. "More user disk space is needed. Better facilities for permanent data storage (such as tape robot systems) are needed."
4. "Will send info later. Must think"
5. "It would be nice with increased computing facilities"
6. "With the current and persistent heavy usage on the two largest facilities, I think the most urgent issue is to fund a major hardware upgrade."

7. "The website (www.notur.org) is in my opinion not very intuitively laid out. Putting up a collection of how-to documents for various tasks and organising them by their relevant sites/machines would have been helpful. If this has already been done somewhere, it could be placed more visibly."
8. "1) The SGI machines have to become much more stable. Having to submit jobs over and over again due to machine crash is not acceptable. 2) Advanced user support should not be "outsourced" as it is now, depending on the geographical location of the applicant. The people who are experts on the platform in question should be given the assignment. 3) Only users who really need a supercomputer should be allowed to use the national resources."
9. "More money for people to help with support and advanced support could be a large benefit for the users."
10. "Be more flexible – PhD and postdocs should be able to apply without having a PI on the application. It should be easier to apply during the whole year."
11. "If possible, be more restrictive to use of forecast-jobs in the batch queue. A general problem with debugging and model development is that it takes long time to get a job started to get it through. Often a quick express job is not enough to obtain the necessary model results. "
12. "Hjelp til parallellisering kunne kanskje vært bedre. Informasjon om denne hjelpen er tilgjengelig."
13. "Get more diverse commercial computational software (not only Gaussian, but eg. MOLCAS)"
14. "Better user support (shorter response times), advanced user support (parallelization and optimization help)"
15. "Perhaps, the total computer resources have been distributed over too many small computers instead of having 1–2 larger installations."
16. "Gridur/embla are a bit outdated, but I guess that replacing these by new computers is a problem of funding which Notur can't do too much about on its own."
17. "Nothing in particular"
18. "More available processors"
19. "Less waiting time in the batch-system. The normal queue is some times slow."
20. "More storage on gridur/embla!!!!!!!!!!!!!!!!!"
21. "Hard to tell. I'm never quite satisfied, but that goes for any user whose work is limited by available computer resources :-)"

22. “More updated development tools such as autoconf/automake/libtool C++ compilers, I/O libraries”
23. “Scrutinize the quality of the staff and the user support at the different centers and look at different ways of using this staff. Perhaps geography should be less important. At least for the application software and user support it seems wise to appoint responsible persons on a national level.”
24. “The main problem (which perhaps is not a failure of NOTUR per se, but rather of the funding situation) is the slow upgrade of the machine. At the current end of the NOTUR period, the demand for computing power is larger than the time available, giving often large wait-times in queue. The services of a non-computer-time providing nature that is part of NOTUR is in my opinion too anonymous. To a very little extent do I see the results of the outreach and TTPs. Although I see the need to market the metacenter structure of NOTUR, I can’t help feel it has made the web site somewhat difficult to navigate in (though it has improved recently). One question is of course how suited the structure and organization is for state-of-the-art benchmark calculations. It appears difficult on all systems to do leading edge, grand-challenge calculation. However, I am personally not convinced this should be a priority area for NOTUR (as I do believe most Norwegian researchers prefer volume over peak), but this should perhaps at least be discussed.”
25. “Savner ressurser for enkeltprosessberegninger. Programvaren på Notur er veldig nyttig for meg, men har ikke noe spesielt behov for parallellprosessering. ”
26. “Have a similar queue system on all supercomputer”
27. “Improve the allocation system, as it is very hard in advance to know exactly how much CPU time you need each half year. Times you are allocated a lot, you may not need much, and then you are given time based on that previous usage. Then the next half year you may need much more than this. ”
28. “Our latest experience tells us that 1) The organization of dedicated user support in order to port and optimize computer code should have more focus on the needs of the users and less on the local educational aspects of the supporters. 2) System upgrades etc. are not OK until the new system works at least as well as the old, FOR THE USERS. The system switch should not be entered into force before this is fully checked.”

### 10.13 Comments on what users think about how NOTUR compares to other such centres

15 comments

1. “no other used yet”
2. “Yes, very good in terms of up-time. I’m only an associated user (primary user is Nils Erland Haugen), so the evaluation from business is a bit excessive.”

3. "Compares very well"
4. "NERSC, USA — they are very, very good at user information, but the system itself is slow and inconvenient."
5. "– Too many small jobs degrading the performance and availability of processor. + Non-bureaucratic staff. (Comparing to NERSC and NCAR in the U.S.A.)"
6. "Auckland University computer Center: NOTUR is much better in nearly every respect compared to how super-computing was run in New Zealand 1998–2000. NOTUR is also much better than the resources that were available to physicists working at the Uni of Oxford ca 1997. They envied us the HPC resources in Norway. "
7. "Not too good when it comes to software upgrades. ECMWF, more defensive upgrade policy, larger staff of course."
8. "Compared to ECMWF Notur is doing quite well (with the available resources)"
9. "I have used NOTUR and its forerunners since 1988, but no other comparable facilities."
10. "ITA Compaq cluster works well but far less users!"
11. "Such comparisons are hardly interesting. The problem is that NOTUR is not a computer center, but rather several, with some added gimmick to make it appear otherwise. The national Meta-Center and the attempts to look like one rather than many are more "Meta" than anything else, at least for the time being. "
12. "I have experience with supercomputer centers in Denmark, Sweden and the USA, and of these different centers, I do find NOTUR one of the very best centers. It is fairly transparent, with in total a large and very helpful support staff. The distributed nature of the project allows you to get detailed help with problems that require more than a discussion over the phone. The centralized allocation of computer time of a rather non-bureaucratic nature allows for easy access to reasonable affluent computer time resources. Clearly, SDSC in the USA is much more suited for grand-challenge kind of applications, but as indicated above, I do in general not think this is what is most wanted by Norwegian scientists using supercomputer resources."
13. "The US military computer network. Their service is good, but their availability is less than NOTUR's."
14. "NOTUR only supports computational Chemistry and comp. geophysics. It is of little interest for people in Scientific computing which develop new methods and techniques. We don't need lots of CPU-hours, but do occasionally need the entire system for scaling experiments. That's impossible (or very difficult) in NOTUR. For these kind of experiments I use and Origin system at TU Dresden and/or a cluster at Univ. of Salzburg. "

15. “Tests of our programs on similar SGI-machines other places (Oak Ridge, USA) shows that EMBLA/GRIDUR are slower for some reason. Also, our programs are faster on gridur than embla (also before the system upgrade).”