



Enabling Grids for E-science

EGEE Collaboration Board questionnaire review

Fotis Karayannis,

Collaboration Board Chair

EGEE All Activity Meeting, July 5th 2005, CERN

www.eu-egee.org




Information Society




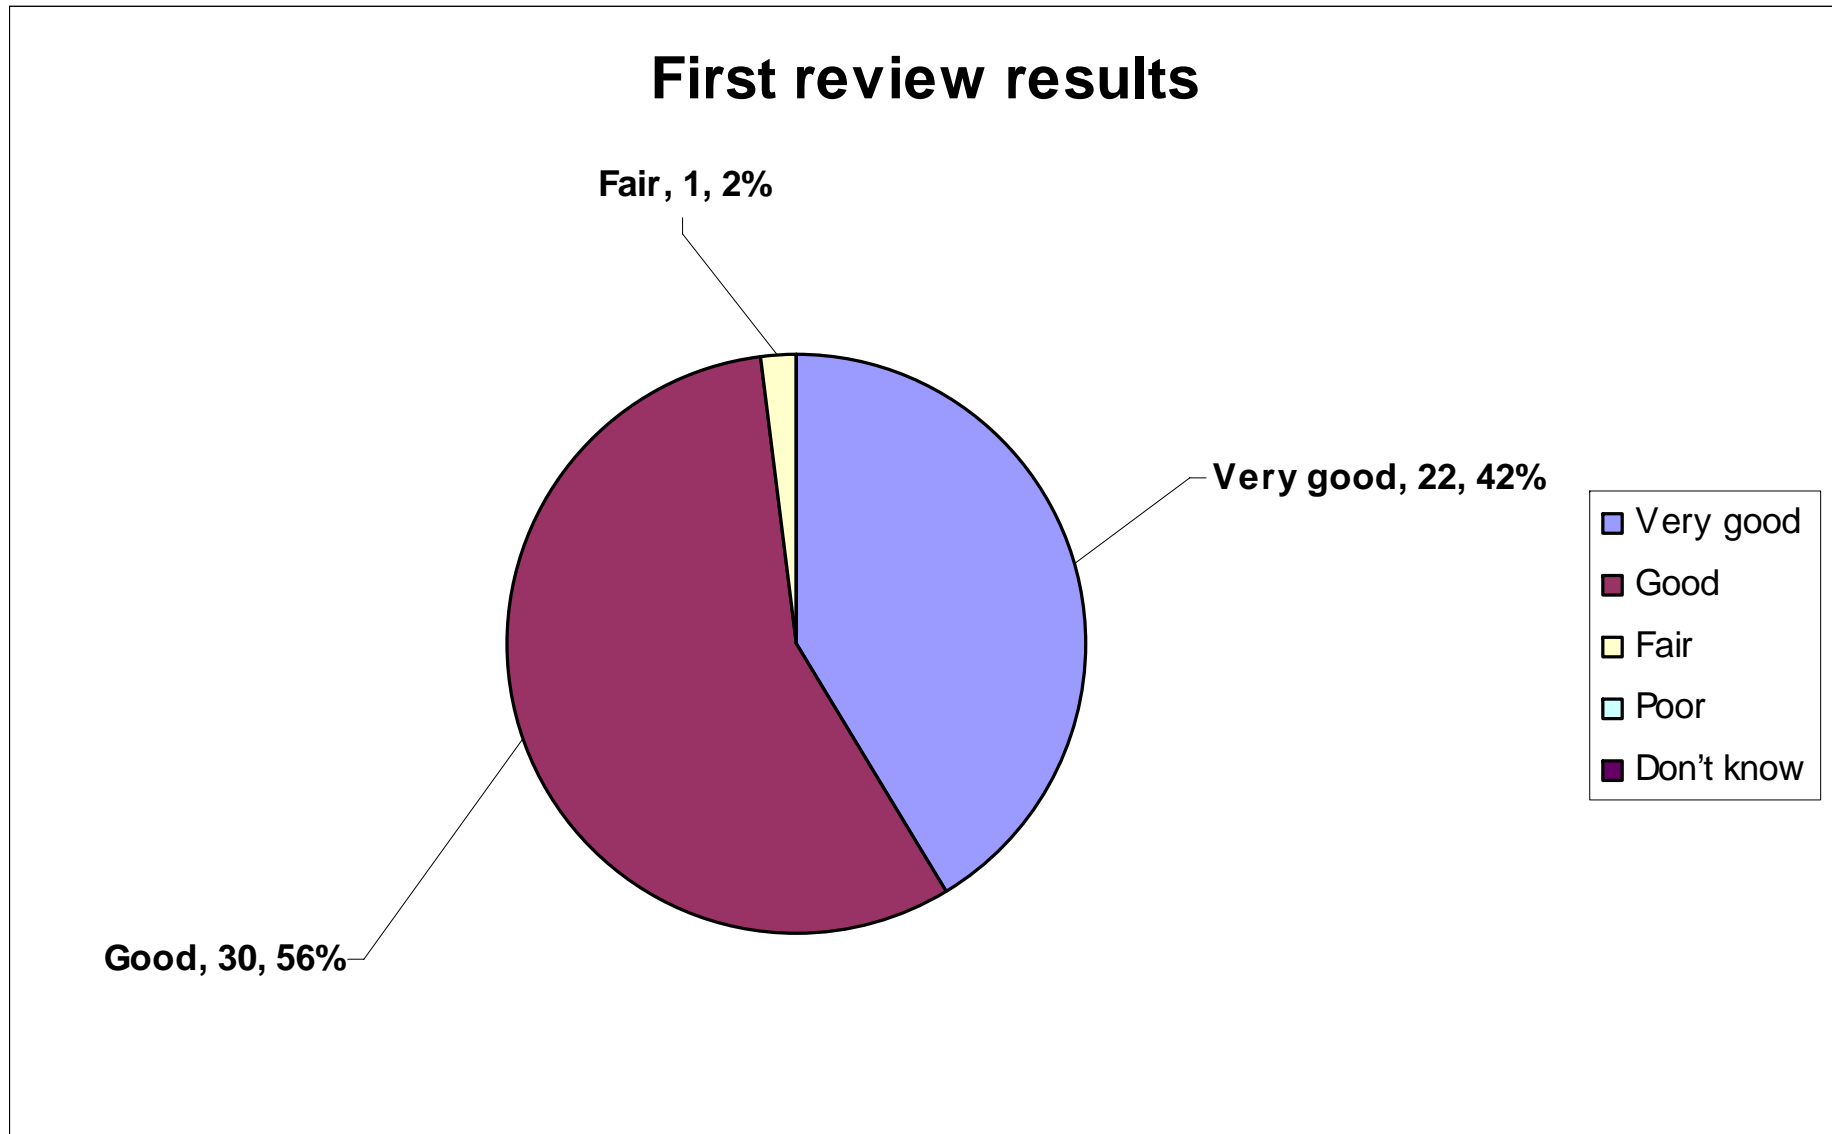
- **Scope:**
 - Feedback from all project partners about the project progress and future directions
- **Procedure:**
 - Agreed during the Athens meeting
 - With a username/pass (per CB partner) i.e. not anonymous (logging in at www.egee-see.org/engine)
- **Type:**
 - On-line questionnaire www.egee-see.org/cb with on-line results
- **Results:**
 - **53** out of **71*** replies → **74,65%** participation
 - *LCG separate vote
 - Results accessible to the Editorial Board
 - Deleted duplicates answers from the Database (last entries kept)
- **All federations answered besides US one**

- **Contact Information**
 - Cross-checking of contact info with user/pass provided
- **General Questions**
 - EGEE progress, Review results
- **Managerial Aspects**
 - PMB, PEB, CB, AFB, EAC, PO, info flow
- **Admin Aspects**
 - Timesheets-Cost Statements, PPT, JRU, Cash Flow
- **Technical**
 - Activities grading
- **EGEE2 related**
 - Minimum changes required, Minimum requirements to join, One partner per country, One per fed, Shift activity focus etc.




Overall progress of EGEE after the completion of 1 project year






- The project has not yet expanded much **beyond HEP**
- The well known **glite delay** influenced a lot in the results
- The intake of **other application groups** is rather limited, esp. if these groups are not large nor as organized as the HEP community.
- 1st year was a real challenge - Well done**
- Glite deployment** should be followed closely by an external group (e.g. based on SA1)
- Most of the milestones for the end of the project have been achieved. The project is widely known and appreciated, as the large number of request for applications are being received.**
- From a technical point of view, the project could be doing better (in particular, regarding the development of **GLite** and its interoperability)
- The project should be **more open** towards new applications and scientific areas. **Quality of sites** must be ensured during registration, not after that.
- Middleware re-engineering** progressing but not in rapid-cycle. **promises of efficiency due to geographical concentration of developers not fulfilled.**
- A1 some very good and some less good (e.g. glite delivery) A2 good .. but focused too much on big pan-European projects and less on really developing an infrastructure


Managerial Summary




The representation of project partners in the PMB

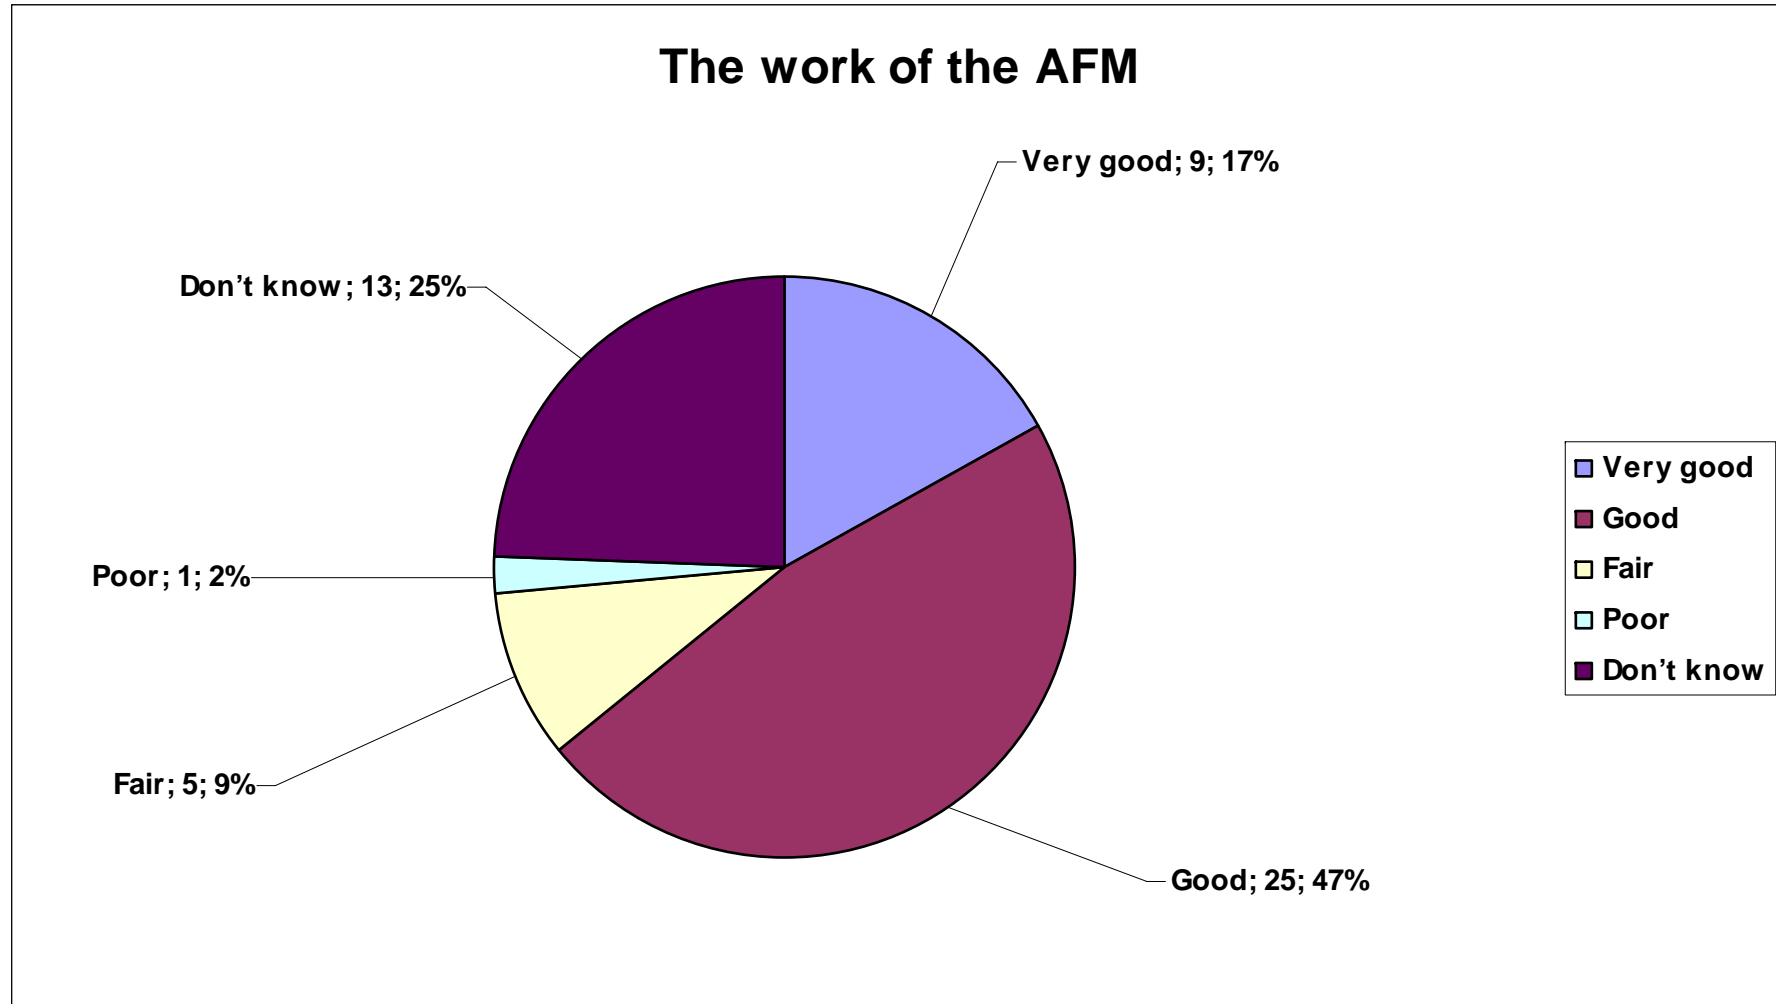


The work of the PEB i.e Activity Leaders




The usefulness of the Collaboration Board







The work of the EAC



The information flow inside the project




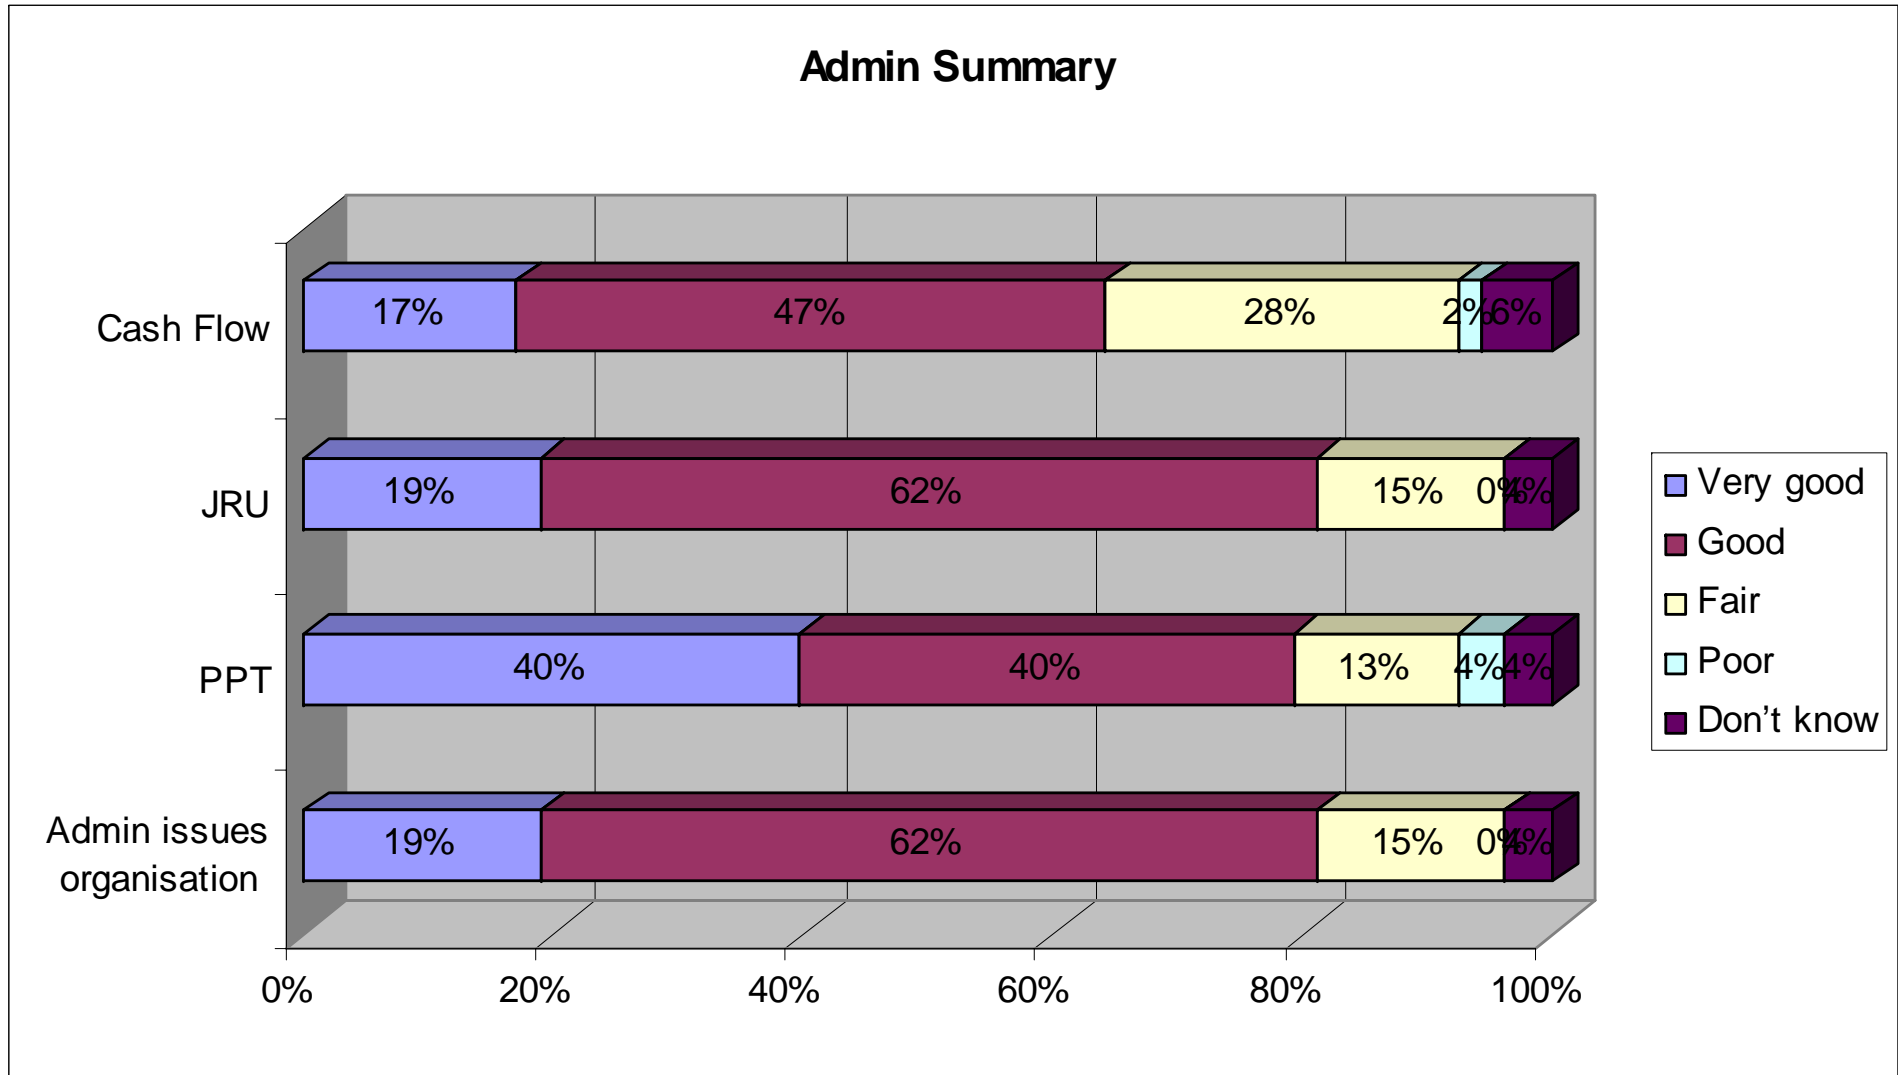
The work of EGEE Project Office

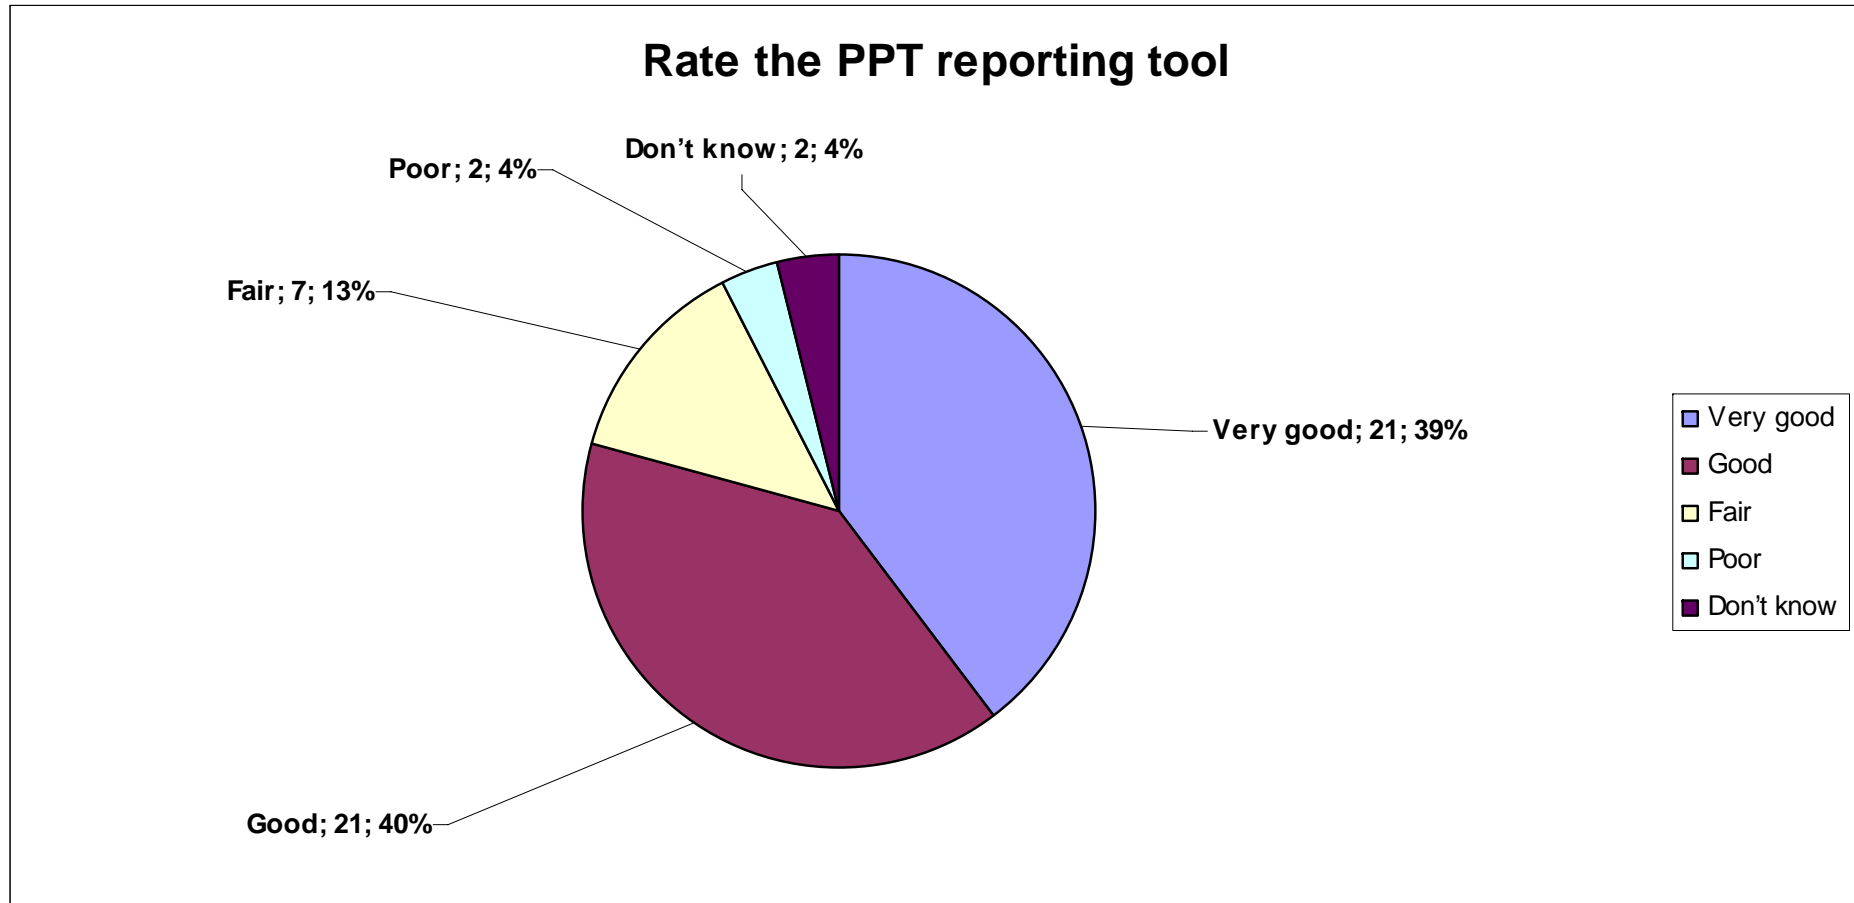


- **Often information arrived in the last minute** before some action had to be taken. I found the **CB** to be useful as far as getting direct information from the PO, however its affect on the project seems to be **negligible**.
- Very successful from the different groups.
- Question 7. Financial matters has been dealt promptly by PO.
- The **IAB and PEB are rather "far away" from the members at large**, but this may not be their fault.
- **-AFM results exceptional** (delays, clarifications, support) - relevant comment from EAC during the preparation of the review - Access to non-technical information is easy. **Technical one is difficult: scattered, sometime duplicated**. Need for a single entry-point (portal)
- It is important that **information flow becomes better** (through regular project updates, newsletters, collaboration board).
- I feel that EGEE partners previously working on the EDG project are very much up to date. Partners coming new to the EGEE project are much less "in the know" and it has taken all this time to gain the appropriate knowledge.
- The **organisational chart works correctly for a large project like this**. Sometimes even the amount of information is so big that is difficult to digest, but this is always preferable. The organisation by federations eases the communication and the vertical organisation, meanwhile the organisation of the activities provides the horizontal organisation. The **EAC has been very important** for preparing the review and guiding the planning of the project.
- The **PEB has spent too long chasing administrative details** and not driving the project forward
- Very often things happen upon too **short notice**.
- **There is a problem with definition and acceptance of real lines of representation, information flow and reporting**


The organization and preparation of the administrative issues in the project

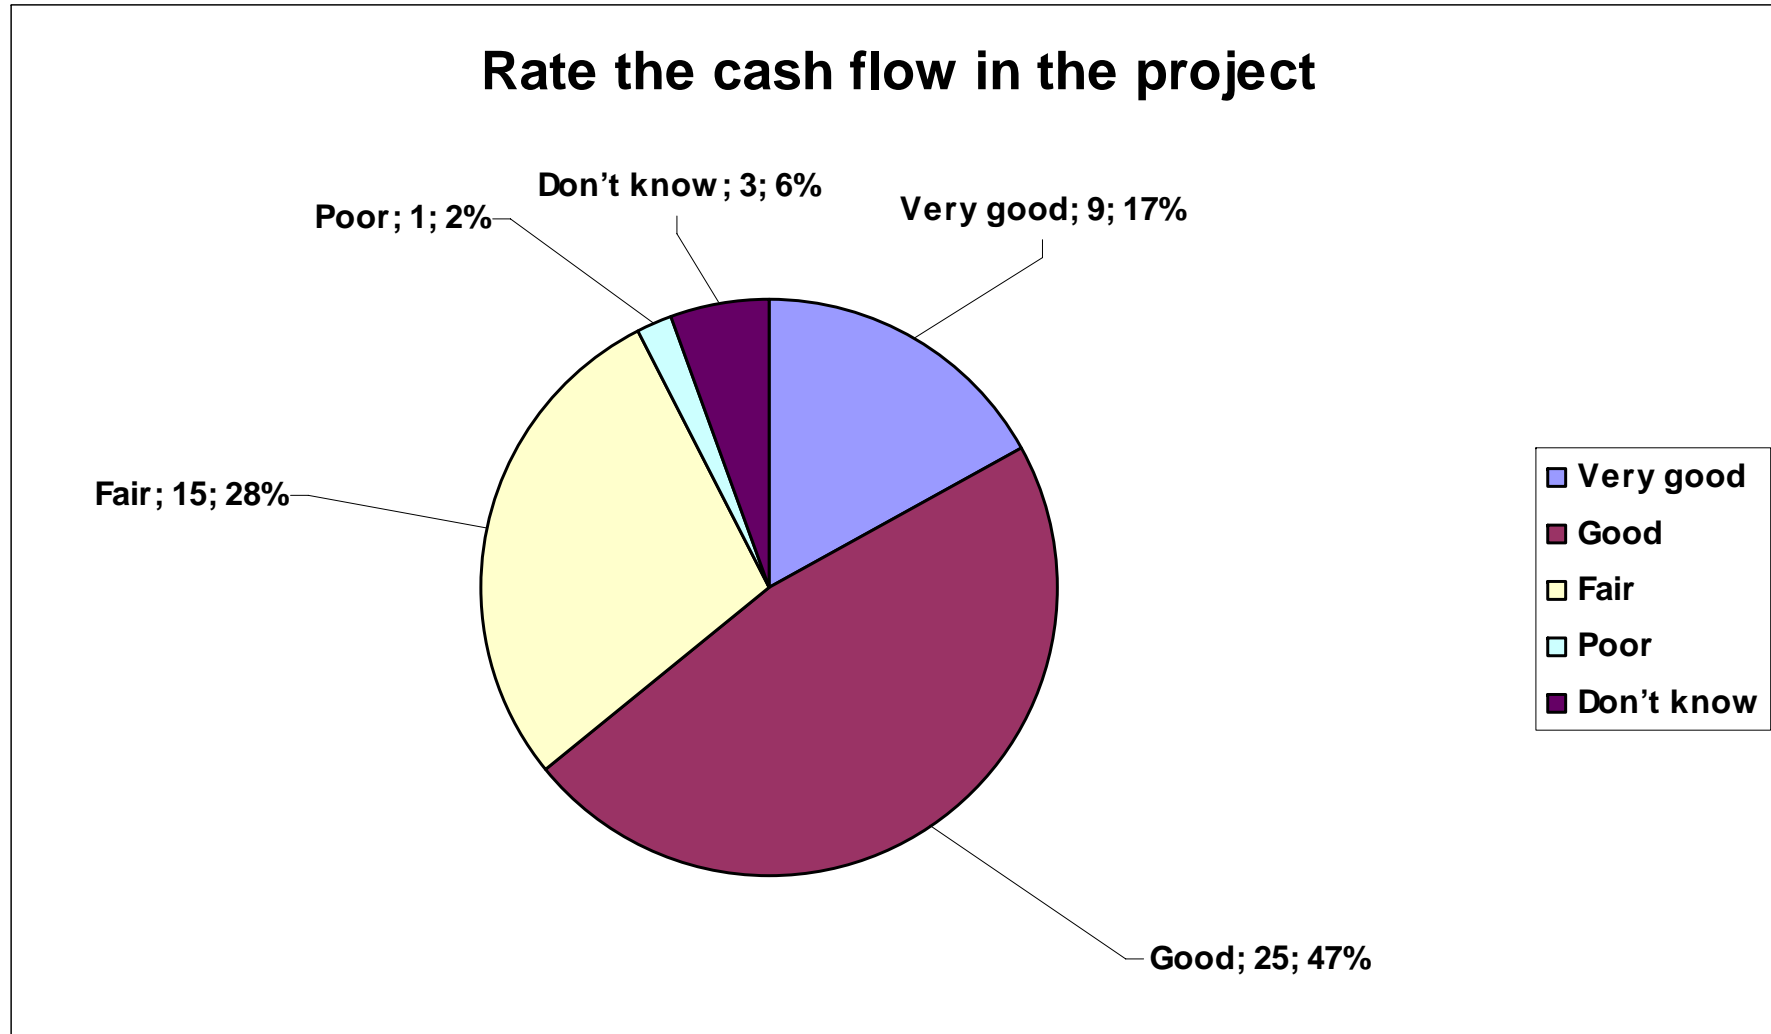







Rate the mechanism of the JRU after one full year of exercise





- My JRU was approved only in February 05 and I still have to **struggle** in order to get it working. *I wish I knew about it for 'one full year of exercise' as question 3 states.*
- Unfortunately, ***one payment was performed with great delay***
- Some ***important questions become clear too late***
- The administrative organisation is **sufficiently effective** to guarantee the correct functioning of the project and the requirements (collecting timesheets, periodic internal submission, audit certificates, etc) are reasonable. Although known, we cannot evaluate the framework in which the JRUs are involved in the project
- We have had problems regarding the **prefinancing** of our activity.
- C1 had a bad start but ended up good
 - *(probably D1 referring to the timesheets)*

Minimum changes are required in the project




Minimum Changes required (2)

- **Justifications**

- since the M/W is changing rapidly the grid is, to a large extent, still an ongoing experiment. Building on acquired experience requires minimal changes.
- Overall, the project has been successful and by making the least amount of changes, we have the opportunity to continue and improve on the work already started
- Depends critically on the separation between application and generic parts of the system
- The project needs a longer than 2 year period to achieve its goals
- **It takes time for such a big project to establish communication network and start with collaboration.**
- But reduction number of partners may help -- **partners with very small contribution may be left out** (or better changed their role in the project, becoming perhaps only affiliates without contract with EGEE)
- The project is people. People are already familiar, connections are already adjusted.
- EGEE-2 should become infrastructure-oriented and application-neutral
- Project **too complex to consider many changes in such a short period.**
- **Concentrate on operation and helping applications** to use what has been deployed in phase 1. Middleware should be only consolidation of basic functionality already developed in phase 1. Leave more advanced middleware development to other projects (timescale too short; incompatible goals between stable and reliable operation and deploying new and untested software.
- Some better matching of skills with activities. Some activities are arranged by region, not according to competence.
- To improve effort on investigation on **MPI based applications**
- **If major changes are necessary, which they are, then major changes should be implemented**
- I think that just minor refinements are needed for an EGEE-2 project, just to maximise the focus of the proposal on the Work programme
- The project has worked well, so minimum changes are required to tune some small problems
- Continuity
- **Distributed ROCs with multiple partners must be rediscussed.**
- Minor changes are necessary to tune the project, but major changes could break a working formula for no obvious gain
- add some from NA4 activities to SA1 (user support)


Set of minimum requirements quantity-wise (CPUs, Storage) for Resource Centers' participation in the project



• *Disclaimer: Due to a DB bug, results based on a small sample*

- 14. For the participation in a Grid project a critical mass is needed.
- 15. Manpower and training are the keys.
- 16. Many operational tasks scale with the number of sites so sites that do not deliver resources should be discouraged unless there are political reasons
- 17. some sites may have more CPUs but small storage, small sites are also valuable


Set of minimum requirements quality-wise (SLAs) is needed for Resource Centers' participation in the project?



Min requirements quality-wise (2/2)

- 1. **Stability and availability** of the systems is important for the users to trust in the services offered by the GRID.
- 2. **Manageability**
- 3. **each ROC can decide on this**
- 4. Once RCs have agreed to share resources, the rules of the game should be established with a SLA.
- 5. At least that they have knowledgeable and willing administrators, but **not full pledged SLA** -- this simply will lead to huge bureaucracy and smaller sites leaving outside (with negativ epolitical impact).
- 6. If any requirements will be applied to RC, then some obligations should undertake GOC, CIC, etc.
- 7. SLA should be signed with the Resource Centers **to support multiple applications and monitor their agreed use percentages per application**
- 8. SLAs are **fine if they are enforceable**. In the short term site validation by daily testing should be accompanied by an effective **black listing approach**. **Sites with insufficient resources or expertise should be removed**. Clearly this requires a review process - and the difficulty will be to avoid the review being politically biased.
- 9. Avoid many small trivial problems.
- 10. Obvious, isn't it ?
- 11. Participated sites must agree to **follow upgrade procedures and acceptable service response**.
- 12. **security reasons, quality of service**

One partner per country should be the ultimate goal for the project?




One partner per country? (2/2)

- **Variety of partners in one country helps to prevent political monopolisation.**
- **This partner should act as the representative of a national grid.**
- Depends on the future user profile in EGEE
- This should only be needed **if the EU Call enforce it**. Actually is not the case
- For any country that has many Grid stakeholders, it is for **their own interest** to get a nationally coordinated Grid body established. One main partner, zero or more JRUs should be the model of choice. **This will limit also the administrative overhead of the project.**
- It's not acceptable for Russian partners because of the financing issues.
- In Russia there are **three independent financial agencies**. However, for the success of the project in Russia the participation of all EGEE partners from related to these financing agencies is necessary. So, one partner per country is not good goal.
- **This model is successful for NRENs and Grids could be seen as NRENs future**
- In Russia no one institution is able to work in the project individually
- **Should be more partners per one country**
- For the manageability of the project a multi-level structure RC-country-federation should be applied.
- **This is the EGEE ultimate vision, following the NREN community model.**
- Only when Grid is like the network will this be possible. I.e. a single representative of national Grid resources. **This is too far off for now** and would exclude necessary participation.
- With the possibility of exceptions if it is clear that one partner cannot represent major interests.
- No, this should depend on the relative skills of each partner.
- We have to trade off between simplicity (i.e. having the federations formally recognized as legal entities somehow) and the need of visibility some (major) partners have
- The complexity of each country must be taken into account by allowing different partners
- **Few partners per country must be the correct answer.**

Do you believe that one partner per federation should be the ultimate goal for the project (with mechanisms like the Economic Interest Group (EEIG))

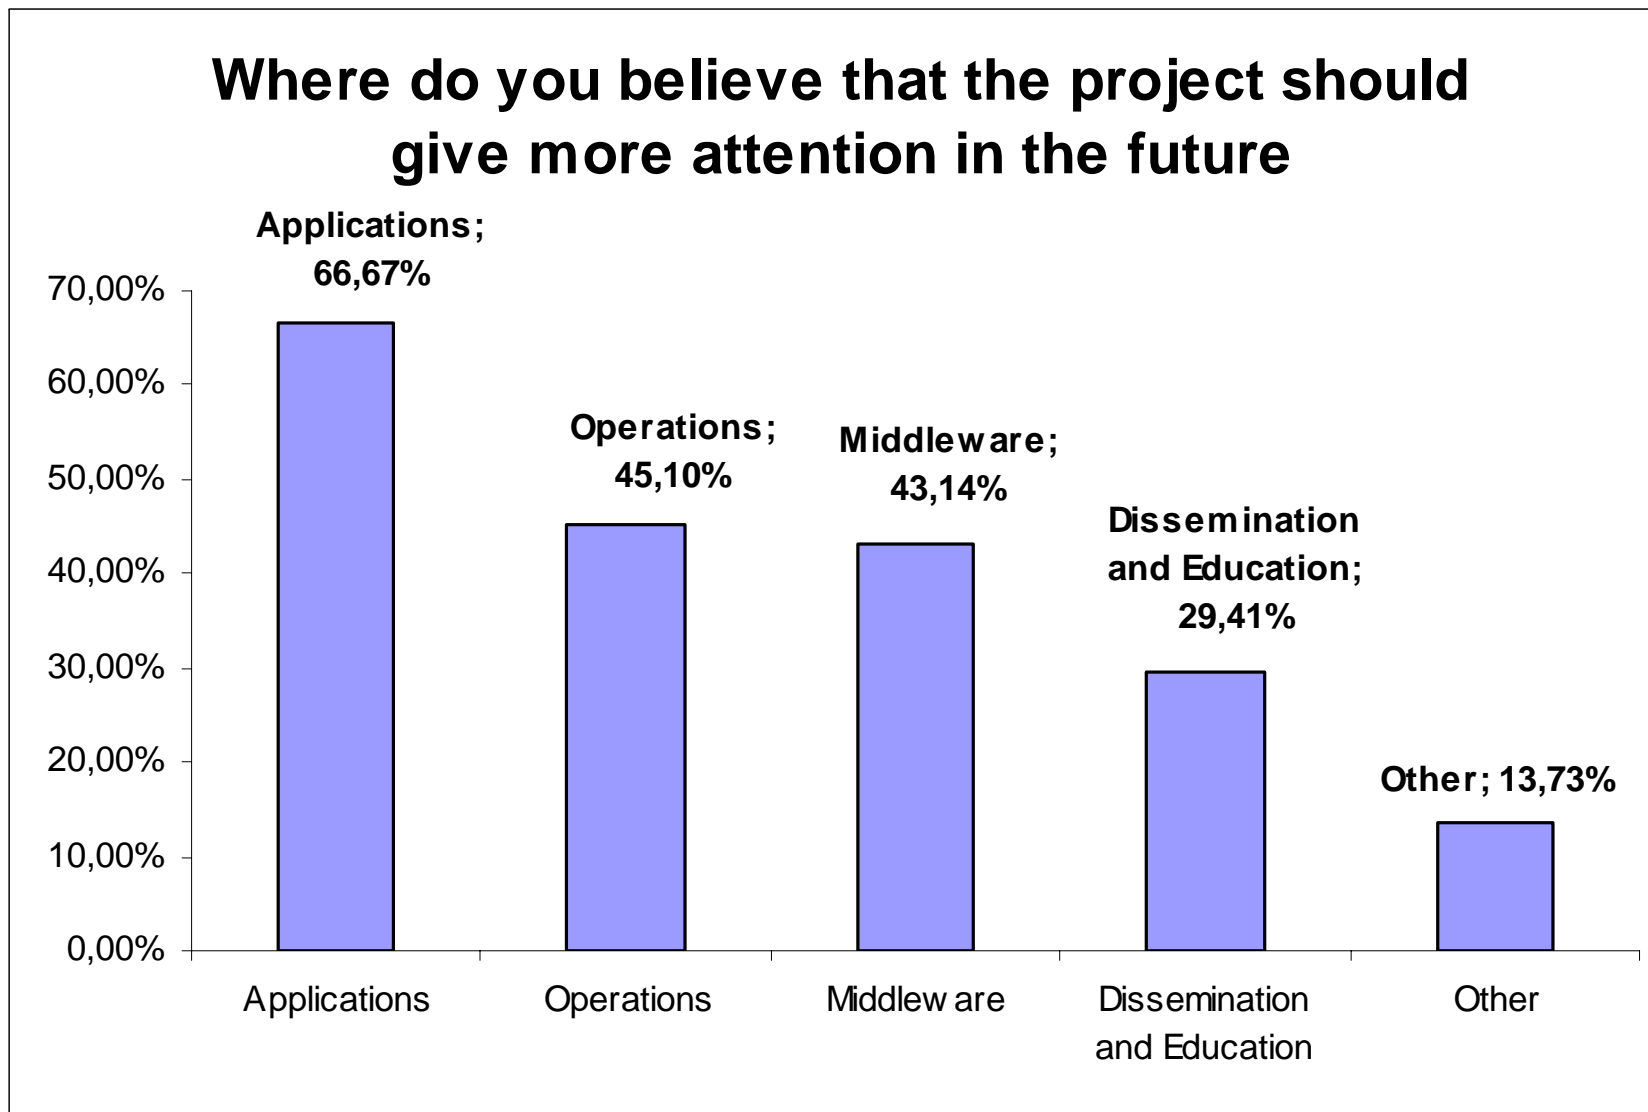
Don't know; 6; 12%



No; 46; 88%



- While the federation is in principle a reasonable way of grouping support services, it is more difficult to see how this works politically. For example some of the federations seem to have distributed their ROC resources across their component countries for political reasons. Better to accept that countries are the basic political unit, with federations used as an operational structure.




- **Other category**

- Marketing - Integration of application challenging as LCG is
- *Industry Forum*
- *Policy and International Collaboration*
- Simplicity
- Consolidating users and attracting *user communities*
- life sciences, industries
- *all* are needed

- **Justifications:**

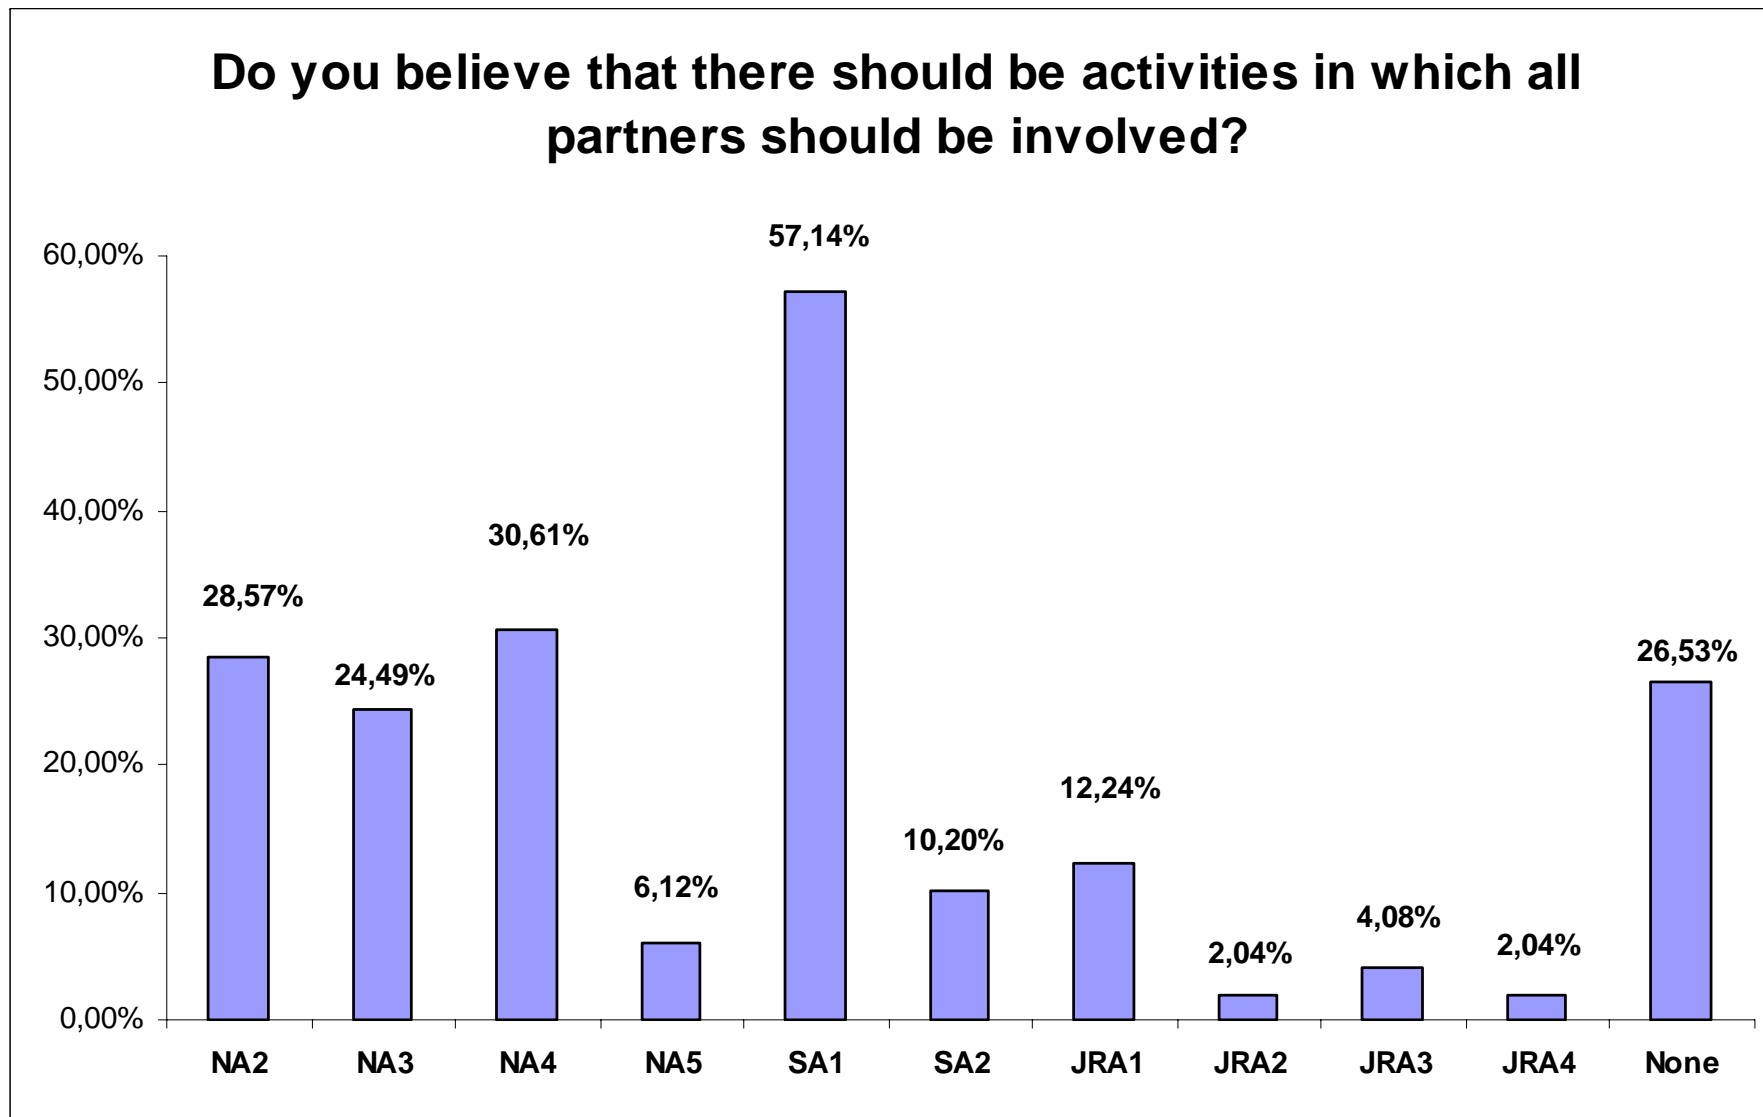
- 1. There are a lot of centers interested becoming RCs but they do **not have enough information**.
- 2. The success of EGEE depends very much on the no of **user groups integrated** at project end
- 3. I think the project has to focus on an **stable infrastructure** and to promote the real use of it
- 4. **Applications** largely will judge the success of the project. Operations, since I have not been convinced by the Federation model of operation.
- 5. The project must make sure that its outcome is not the general feeling among scientists that Grids are just for large application groups. Project must help to bring more groups, esp .smaller ones (for which e.g. **training** at the other side of Europe is too expensive to participate).
- 6. Operations were the main focus of EGEE. More focus on Applications and Dissemination/Training. Policy issues are also becoming very important, as well as international collaboration with satellite and national projects

- **More justifications:**
 - 7. The goal of the project is an **operational grid** for multiple sciences. It is very important to attract KEY scientists by providing a good, reliable service and helping them to use it.
 - 8. **Applications** are what the users ultimately need to use and these really drive the middleware requirements.
 - 9. We should aim at an equal share between **Operations, Development and Applications**.
 - 10. You need **applications** to run on an infrastructure otherwise its useless. You need strong expertise and new tools to bring applications on the GRID.
 - 11. to have satisfied **users**



- **Justifications:**


- As in 6, **all** the activities are of great impact in the organisation of the project . It will be difficult to maintain the well structure and organisation if the activities are not kept
- **JRA4** type of tasks should be done in **GN2**
- Continuity is crucial
- For **JRA** activities, as well for **NA5** we do **not** have any **clear opinion**.
- **NA3 and NA4 should be connected**, as well as **JRA1 should "consume" the JRA3 and JRA4**. More tight collaboration between JRA1 and SA1 (on daily basis, probably leaving part of current JRA1/3/4 activities to SA1) is needed. JRA2 should work more closely with JRA1 and SA1
- I do not know really. It may depends on other project progress.
- Don't know much about unmarked activities
- All of them are needed to arrive at a truly operational Grid service.
- JRA3 and JRA4 should be integrated in other activities
- Others should be integrated into the main JRA1 and SA1 activities.
- International collaboration has been very well done in phase 1 - consolidation on the core is needed now. Only lightweight liaison with research networks is needed - GEANT and the NRENs are mature organisations. High bandwidth networking will be brought in automatically by the VOs that need it (just as it is the VOs that bring in the resource centre capacity). Network research should be left to other projects. JRA2 and JRA3 should be embedded in SA1 and JRA1. JRA1 should be reduced in phase 2 to providing basic grid functionality, more a software distribution than a development activity.
- The project could benefit from re-visiting the overall structure in the light of what has been achieved. New goals should be set, not a continuation of the old ones.
- This is an operations project. Must keep internal development as low as possible.
- **Testbed and user support**, but all others are also needed

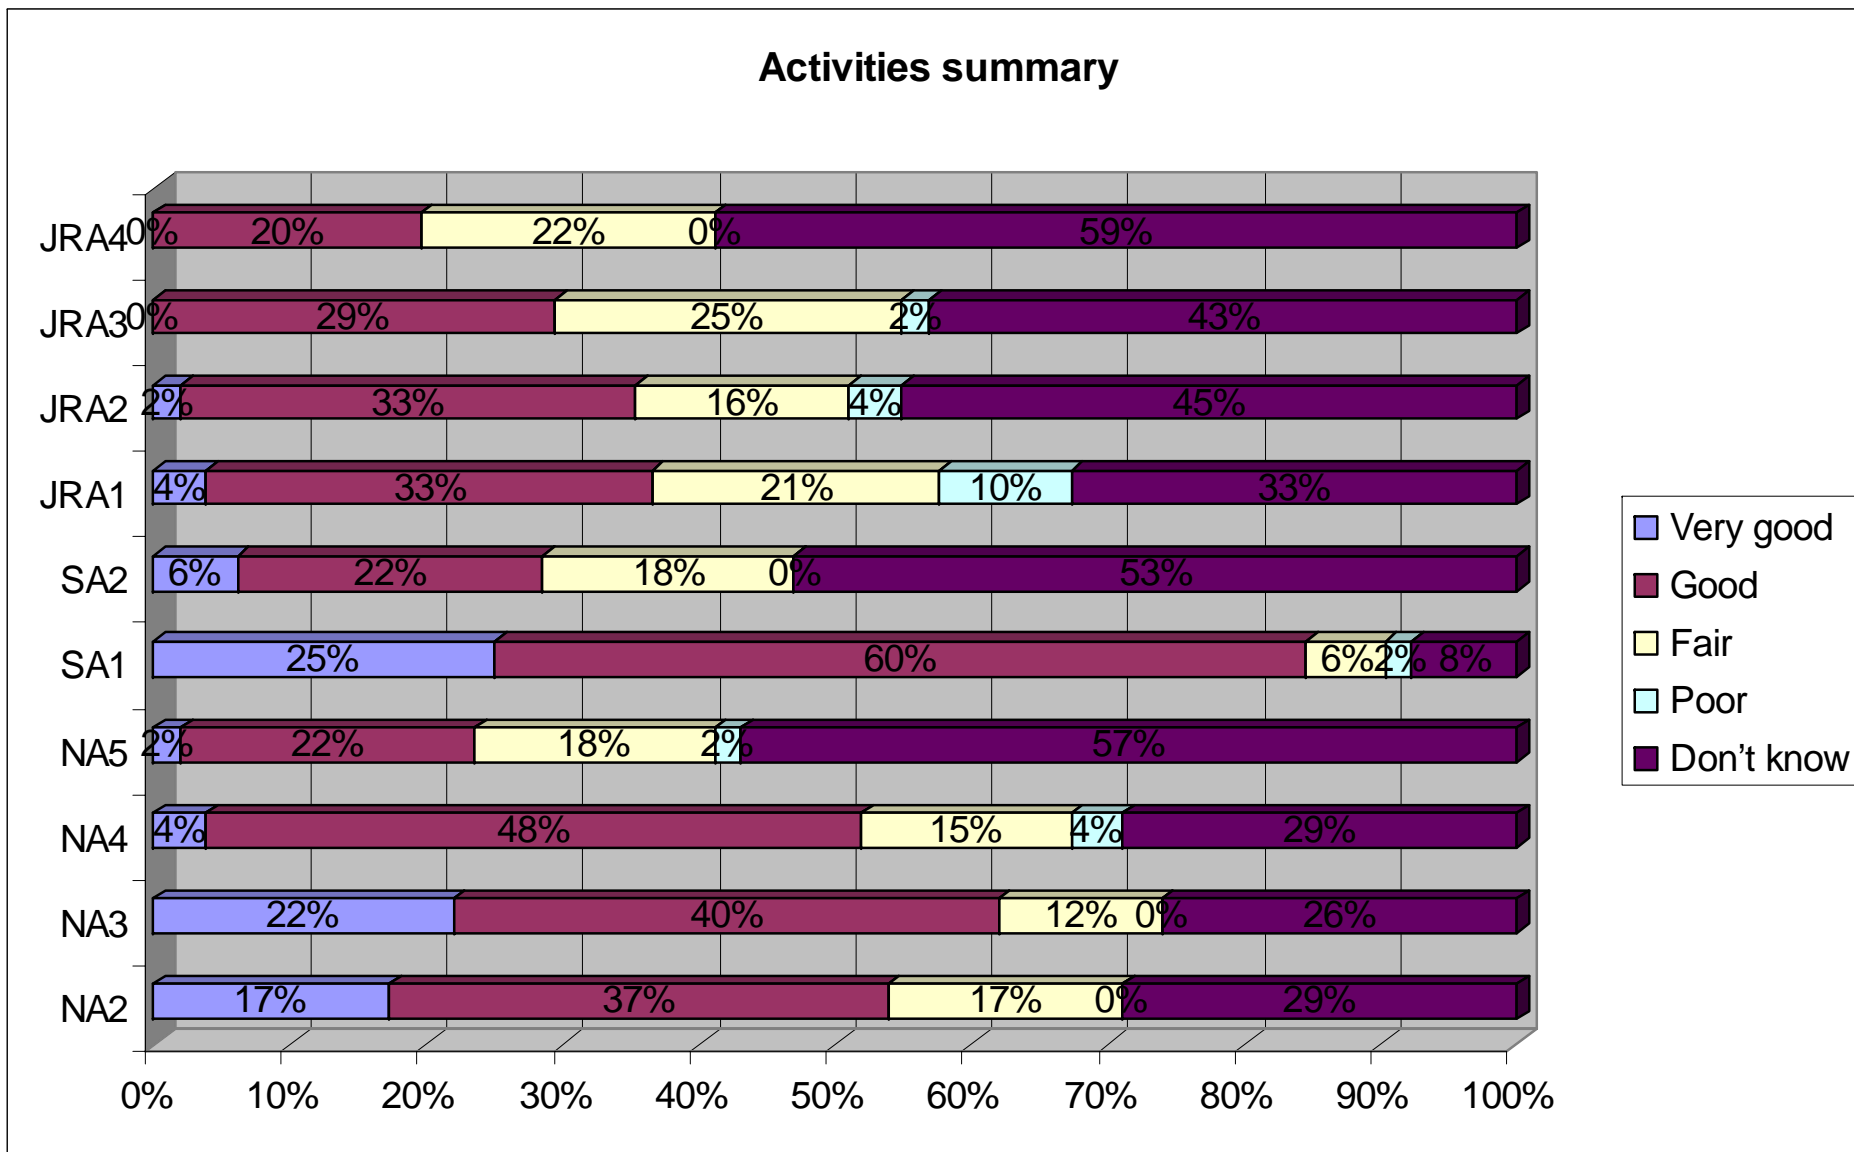


- **Justifications:**


- all should have grid sites, hence involved in SA1, and all should expand the the number of VOs and fields of interest
- Don't see a justification for this.
- I think it depends on the partner.
- These are activities that benefit from large geographical coverage.
- The good JRA1 site may not be also involved directly in any other activities. The same applies to good NA3/4 sites, although association with SA1 would definitively help.
- Each partner has its own force and weaknesses it is difficult to do everything in such a big project.
- **Least one partner/country in NA2 & NA3**
- **On federation level.**
- I cant see that this would benefit anyone.
- The basis of the project is the exploitation of GRID technologies and I think that **all FEDERATIONS should be involved in JRA1, NA5 and NA4**
- **A working Grid infrastructure is the basis. Everybody should participate to learn how it works.**
- Against common sense
- SA1 combined with at least one more activity

Please rate the performance of NA2 activity







Please rate the performance of NA3 activity




Please rate the performance of NA4 activity




Please rate the performance of NA5 activity




Please rate the performance of SA1 activity



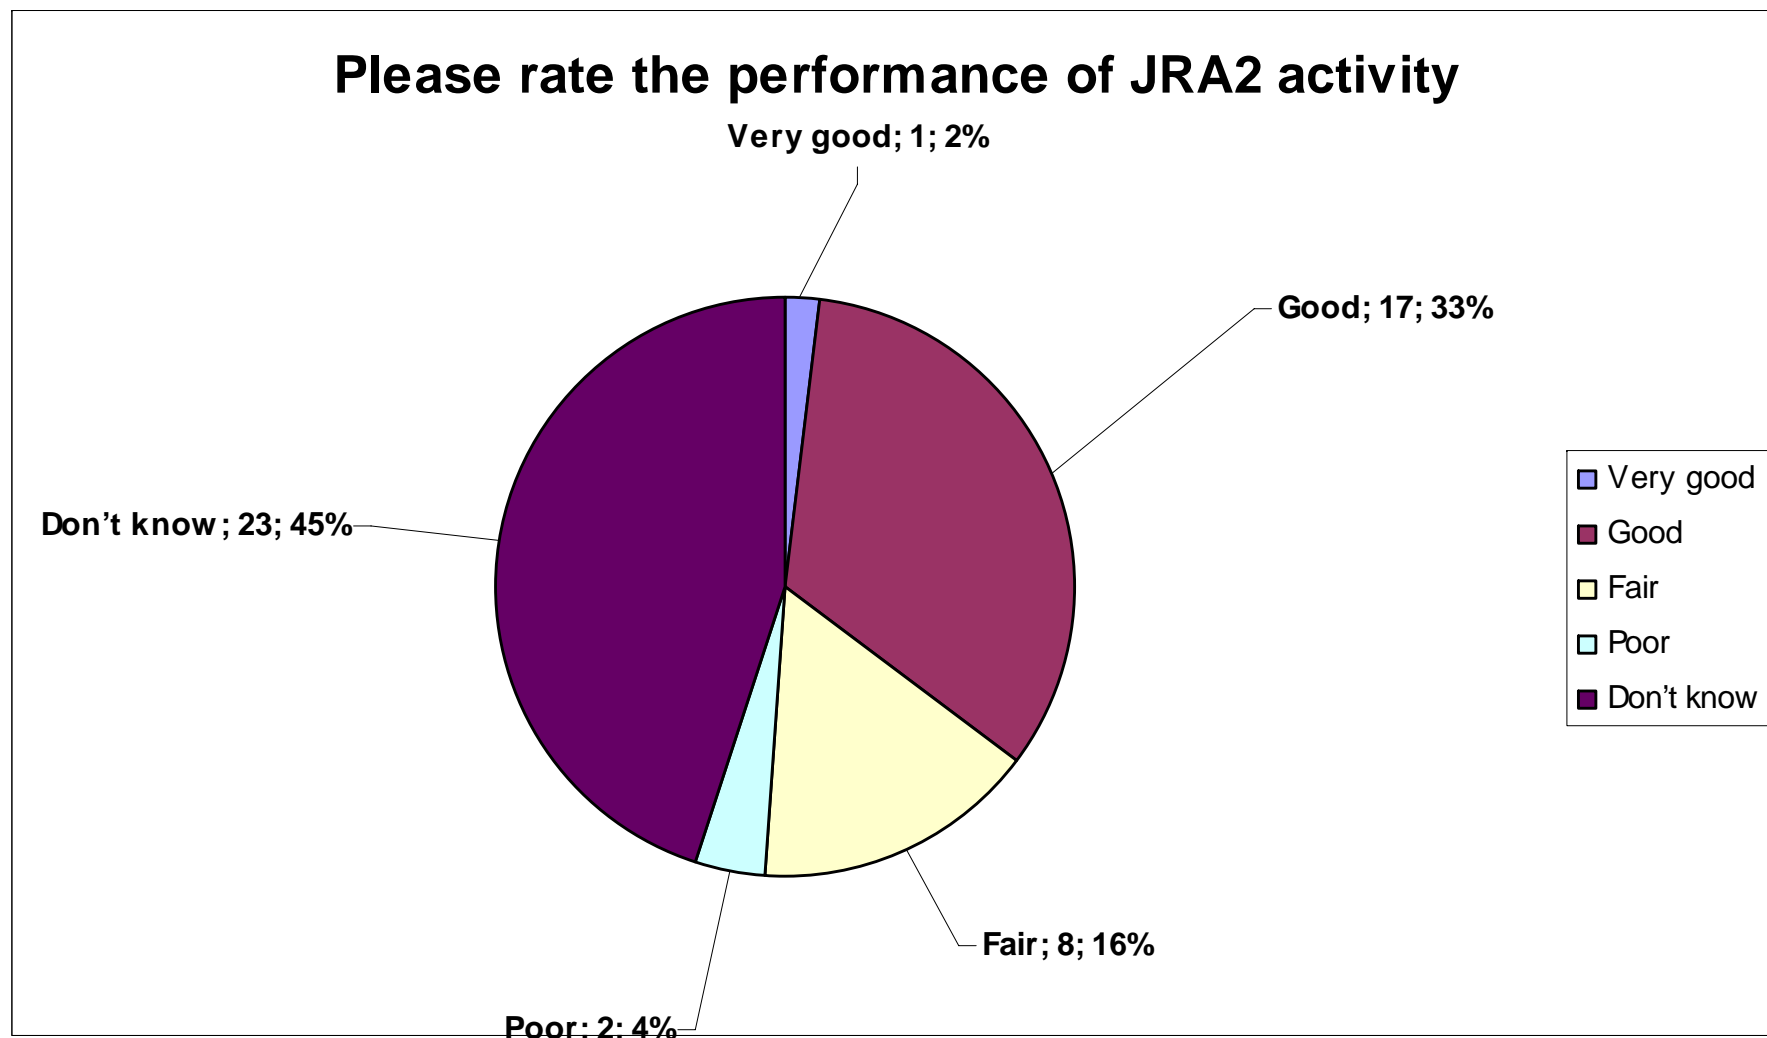
Please rate the performance of SA2 activity




Please rate the performance of JRA1 activity




Please rate the performance of JRA2 activity



Please rate the performance of JRA3 activity



Please rate the performance of JRA4 activity



For activities in which we are not involve, it is a bit difficult to rate the performance. However, on the conferences the Technical Presentation of the different activities have been very good.

It is near to impossible to evaluate the activities we are not involved in. The "Fair" for NA4 reflects evaluation on the predominant interest in large application groups and very limited help to small groups (but still important, as there is many of such groups)

Lack of 64-bit OS support gLite still not deployed

-gLite deployment not significant - **Need for transverse security policies**, not only local site dependant. Each incident should be analyzed and diagnosis and action plan should be sent to the whole community

NA2 should have more specialised persons, NA3 does not seem to be a live group of people (e.g. like SA1), NA5 should activate more partners, JRA1 needs external monitoring (e.g. from an SA1 team),

I do not think it is appropriate to ask such questions

The dissemination (NA2) could be improved with a tighter connection to the local partners. The training (NA3) has been very extensive, although some courses could have needed a consolidation on the subjects treated. NA4 will improve with a wider involvement of user communities. SA1 should concentrate more on reliability and robustness for production use. JRA1 has done a good job considering the timeline restrictions, although some requirements are still missing. More visibility of the activities could be needed to leverage the usage of their results (JRA3). In the case of the JRA4 and SA2 (Network Activities), although known, we have not the sufficient skills to evaluate the quality of the actions. Finally, in the case of the NA5, We do not have enough information to provide a reasonable evaluation.

Output of production quality M/W is too slow to attract developers and users during the project lifetime. The virtuous cycle hence has not really started to work.

Support for bringing in non-LHC HEP applications is poor. In particular the procedure is not clear.

1. EGEE-2 should complete the gLite middleware plus to continue efforts aiming to creation the veritable pan-european grid - biggest, secure, reliable etc.
2. **A more clear connection between SA1 and other activities should be considered in EGEE2, allowing participation of SA1 people in other activities. One clear example is what happened with glite and training.**
3. no comment.
4. **Regional Operations Centers and PMB members should hear more about their Federation members.** My country was missing from Federation mailing list and phone meetings practically whole year regardless of MoU. I would thus prefer that each country has a presence in decision making boards to only federation representatives.
5. **More application orientation. Better inter-activity collaboration.**
6. **Need of : - a new great challenge - strong dissemination activities involving partners in every country,** including marketing aspects - very high service level agreement - middleware leadership - infrastructure and middleware sustainability
7. **- EGEE II should be closer to Applications and improve the interface between App requirements and M/W** (today we are neither motivated nor confident on this) - Explore the possibility to sub-contract gLite maintenance to a commercial company (start-up - a new "Red Hat") to create interest in the Industry
8. It is a huge and complex project and in general the progress was quite good
9. The project is evolving positively. **The integration of the activities of potential new projects which have been submitted recently will be an important issue.** The coexistence of internal and external activities will be a major challenge.
10. We propose to establish workpackages dealing with complex workflows and data management.
11. It is astonishing that the access control to this questionnaire was not deployed via EGEE certificates.

- **Thanks!**
- **For more information:**
- <http://agenda.cern.ch/askArchive.php?base=agenda&categ=a045529&id=a045529s0t4/moreinfo>
(excel sheet)