



The e-IRG Roadmap 2010 public consultation Comment summary and e-IRG board response [April 2010]

This document contains complementary material to the e-IRG Roadmap, published on 8th April 2010 at the e-IRG website (<http://www.e-irg.eu/publications/roadmap.html>).

www.e-irg.eu

e-IRG
e-Infrastructure
Reflection Group

e-IRG secretariat
P.O. Box 405
FI-02101 Espoo, Finland
Phone: +358(0)9 457 2281
Fax: +358 (0)9 457 2302
secretariat@e-irg.eu

Visiting address
Life Science Center,
Keilaranta 14, Keilaniemi
Espoo, Finland

Foreword

The development of the common research e-Infrastructure has been very rapid since the publication of the latest e-IRG Roadmap update in 2007. The numbers of users, organisations involved and research activities supported have grown dramatically. However, this pales in comparison to what is waiting for us just around the corner: large-scale e-Infrastructure initiatives that are supporting activities ranging from major global leading-edge endeavours to new kind of discovery processes of individual scientists.

The research e-Infrastructure needs to act as a substrate that supports a wide range of innovations stemming from new, sometimes unplanned and spontaneous inter-disciplinary activities the common e-Infrastructure has enabled. Due to these ongoing changes, it was seen that a completely new structure and approach for the roadmap document was needed to capture the vision of the future. It was also clear from the start that a forward-looking policy document like the e-IRG Roadmap cannot be produced in isolation.

The expertise of the e-IRG delegates creates a strong foundation on the issues related to the e-Infrastructure and intricacies related to producing reliable and cost-effective services for the expanding user base. This expertise was essential to make sure that new usage scenarios were adequately taken into account when formulating the policy recommendations. The e-IRG board also thanks the ESFRI appointed experts that provided continuous valuable support for the creation of the e-IRG Roadmap document, starting from the early draft versions all the way to the version that entered public consultation in December 2009.

The e-IRG board would equally like to thank everyone who participated in the public consultation process and contributed their time and expertise to the process. The comments received allowed further refinement of the document. We also think they confirmed, to a large extent, the broad applicability of the key messages and recommendations of the Roadmap document.

The present document constitutes a genuine attempt to leverage this input beyond the Roadmap itself. Especially when it comes to the ability of the e-Infrastructure-related policies to support fortuitous discovery and innovation, establishing and maintaining a dialogue with a large group of stakeholders on the global scale can be as important as the e-Infrastructure policies themselves. Common policies and practices are necessary to make the e-Infrastructure services available to and usable by the user communities, but these policies and practices need to evolve along with the needs and practices of the users. As a step to maintain this dialogue and evolution, this document not only publishes the comments themselves, but also uses the comments as a way to illustrate the motivation and background of the editorial decisions that were made.

Our sincere hope is that this document, together with the e-IRG Roadmap itself, will support open exchange of ideas and best practices in the e-IRG workshops and other participatory e-IRG activities in the future!

On behalf of the e-IRG board,



Leif Laaksonen
e-IRG chair

Summary of e-IRG roadmap comments and e-IRG responses

The following table summarises the results of the public consultation of the e-IRG roadmap (between 21st December 2009 and 31st January 2010). This document can be downloaded from the e-IRG website, from the same page as the e-IRG roadmap itself (<http://www.e-irg.eu/roadmap/>).

From	Comment (paraphrased)	Action
Philippe Gavillet (comment 1, page 10)	Avoid the impression that e-Infrastructure drives re-organisation of research groups	Response: e-IRG sees e-Infrastructure as a domain where various needs for organisational change can become visible and culminate, since e Infrastructure is a key enabling tool for collaboration between individual researchers and research groups. However, e Infrastructure should not force re-organisation of research groups, at least beyond possible technology-derived issues that may make certain collaboration models easier to support than others.
Philippe Gavillet (comment 1, page 10)	Emphasise that outsourcing of CPU and data-intensive tasks to global e-Infrastructures should not weaken the necessary local resources. Consider downsizing e-Infrastructure to run locally as standalone solution	Response: the role of the global e-Infrastructure should be to empower local support organisations to work more efficiently by a) relieving them of some routine tasks and b) establishing links that allow knowledge sharing. Local support is crucial, but its role can be developed – e.g. by concentrating efforts on more advanced, customer-specific support tasks. Standalone solutions are somewhat orthogonal to this development. While e-IRG recognises that in some cases the ability to run e Infrastructure services autonomously is of interest (e.g. in cases where network outages are to be expected), this model should be considered mainly a backup option.
Carlos M. Romeo-Casabona (comment 2, page 10)	Perhaps rethink recommendation 3.8 from a more realistic point of view: it might not be easy to extend the provision of resources for some user communities in new EU member States.	Response: there may indeed be challenges in equal resource provision across the whole group of EU member and associated states. But the goal should be that e Infrastructure reduces the differences between the states. To address this issue, recommendation 3.8 suggests finding methods to help accomplish this.

<p>Prof Alex Reid (comment 3, page 10)</p>	<p>Emphasise the point that researchers need capacities that exceed commercial services. We agree that commoditisation is a goal worth pursuing, but doubt it can be reached in the short term (especially in networking). Based on 20 years of experience in the governance and financing model of AARNet (member owned, operated and funded), we recommend volume-based charging of third party use in ad hoc cases, and subscription-based charging for sustained collaborations.</p>	<p>Response: e-IRG finds AARNet’s experiences very interesting and hopes there will be further opportunities for dialogue between AARNet and European e Infrastructure practitioners. Regarding (commercial) commodity services, it is obvious that the move to commodity services should not be pursued as a dogma. However, it is important to have processes in place that facilitate cost/benefit comparisons between academic and commercial solutions in cases where their technical capacities and capabilities are similar.</p>
<p>Dietmar Erwin (comment 4.1, page 11)</p>	<p>Include a short, clear definition of “e-Infrastructure”</p>	<p>Roadmap edit: a definition is now included in the part “e-IRG and its mission”, including the list of current components.</p>
<p>Dietmar Erwin (comments 4.2, page 11 4.3, page 11 4.6, page 12 4.7 page 12)</p>	<p>Improve readability of the text, correct the bullet list in the executive summary to match the recommendations, readability and consistency improvements</p>	<p>Roadmap edit: e-IRG would like to express its gratitude for the suggestions for improvements that the separately delivered document contained. Many of these suggestions resulted in clear improvements to the final version of the roadmap</p> <p>Response: Regarding mention of PRACE in the HPC recommendation: the text of the recommendations was refined at the e-IRG meeting of November 2009, with the goal of keeping the length of the recommendations as short as possible.</p>

Dietmar Erwin (comment 4.4, page 11)	Clarify the concept of commodity computing	<p>Roadmap edit: a footnote explaining the broad definition used in the document has been added to 2.4.3.</p> <p>Response: it is obviously a challenge to provide a concise definition, especially where there is blurring between several traditional and emerging technical – and marketing – categories in this area of computing. Defining commodity computing as “anything but HPC” might be possible, but not much more precise than the current informal definition.</p>
Dietmar Erwin (comment 4.5, page 12)	Define the role of the roadmap and roadmap recommendations in the e-IRG process.	<p>Roadmap edit: a minor step in creating the link is referring to recommendations as roadmap implementation steps on the heading level.</p> <p>Response: however, the e-IRG’s approach to execution of its mission is perhaps less centred around the major documents than the e-IRG website, for example, might lead one to believe.</p>
Dietmar Erwin (comment 4.8, page 12)	Clarify assumptions on data production: I find it hard to believe that the flood of published knowledge in 2020 might rival the flood of raw data today	<p>Roadmap edit: a minor clarification was added to explain that this was partially based on assumptions that the methods used to present research results and the distilled knowledge may also increase the amount of data contained in a single publication by orders of magnitude.</p> <p>Response: presenting meta-analyses in a way that includes derived data sets (together with necessary means for verifying their origins and methods used to produce them) can also play a role.</p>
Thomas Eickermann (comment 5, page 12)	Acknowledge that HPC has been acting in the service provision model on national level for a long time; move Exascale project reference to first mention of the project	<p>Roadmap edit: the European scope of the service provision added to 2.2; Exascale reference corrected (2.3.2, 2.3.3).</p>

Andrew Lyall (comment 6, page 12)	Take into account the data-related requirements of the ESFRI Biology Research Infrastructures; establish a permanent body ensuring that relevant stakeholders meet regularly.	Response: e-IRG acknowledges the importance of biology research for Europe and sees it as one of the very important e-Infrastructure user communities. The increasing collaboration between e IRG, ESFRI and other major stakeholder groups creates a shared, dynamic structure that is well placed to ensure that sufficient information exchange happens between e-Infrastructure service providers and researchers in the biology domain.
Peter Wittenburg (comment 7.1, page 13 and 7.5, page 14)	Mention the importance of trust; commercial solutions often do not meet the requirements of users (and even if they do, trust is an issue)	Response: trust is indeed very important, but beyond quantitatively measurable metrics (service level agreements, etc) the prerequisite for trust is creating links between parties that should trust each other. In the case of commercial entities, establishing and maintaining this trust may take more time and effort. It may even require companies to relocate into countries where the legal framework governing privacy of electronic information is sufficiently advanced and expected to remain stable.
Peter Wittenburg (comment 7.2, page 13 and 7.6, page 14)	Address service oriented architecture (SOA) and the general sustainability model for software.	Response: the exact definition of service-oriented architecture is at this time somewhat open. However, if SOA fulfils its promises, it should be possible to integrate it seamlessly as part of the implementation technologies of e Infrastructure. The eventual link between e Infrastructure and SOA seems to be most natural in the commodity computing area. The sustainability models for application software might emerge more naturally if they were linked with the archives of the research data itself (on a per research group/collaboration basis). This would also provide a natural interface to the common e Infrastructure via data-related services.

<p>Peter Wittenburg (comment 7.3, page 14 and 7.4, page 14)</p>	<p>Consider the stability of minds and societies to be a grand challenge; domain-specific knowledge is needed to build suitable interfaces/services.</p>	<p>Response: e Infrastructure can only play a role in addressing any grand challenge through the efforts of its users. In some cases the impact of user input can be unanticipated (as illustrated by the history of the World Wide Web), and in other cases, quite predictable (as for certain parts of “traditional” RI use). Perhaps studying global research collaborations and their innovation processes might help in developing metrics that would allow monitoring of the “social climate change” mentioned? However, this kind of research would require an approach that would not be limited to e Infrastructure – although e Infrastructure could offer efficient tools for gathering data.</p>
<p>Peter Wittenburg (comment 7.7, page 14)</p>	<p>Consider which comes first: standards or interoperability?</p>	<p>Response: there are some notable examples of cases where standardisation proceeded well in advance of the development and testing of interoperable solutions (especially interoperable solutions from different sources). This tended to lead to delays and required extra effort for eventual implementation. For this reason, the roadmap wanted to emphasise the importance of practical interoperability testing as part of the standardisation strategy.</p>
<p>Mirco Mazzucato (Comment 8, page 15)</p>	<p>Consider that the coupling of different roles among early adopters has supported innovation, while not preventing the move towards the service-oriented model.</p>	<p>Response: initial experimentation within a fairly well organised community can indeed result in relatively free service operation and innovation, even when using fundamental components and interfaces. However, as the user base grows and diversifies, maintaining this innovation potential on a larger scale and for more diverse user groups will likely require different approach.</p>

Edward Aronovich (comment 9, page 15)	Complement IaaS with PaaS and AaaS models.	Response: expanding general purpose “component libraries” to higher abstraction levels is a positive development. However, the most efficient way to achieve this may be through generalising components that can be used in several domains from successful (usually domain-specific) applications. Another issue is the fluidity of these concepts (IaaS, PaaS, AaaS): the terms are in some cases used almost interchangeably, making them problematic when communicating the features of the e Infrastructure to new users and user communities.
Kees Neggens (comment 10, page 15)	Consider the need for interoperability of organisational and financial models	Response: The issue of transparency and interoperability in models for e Infrastructure service provision is indeed an important one, and more resources should be allocated to solving this question. Roadmap edit: the executive summary has been edited to better reflect this issue as presented in section 2.4.3 and to highlight it. In addition to this, the related recommendation 3.2 has been expanded to highlight the issue (discussed with all e-IRG delegates)

<p>Christos Kanellopoulos (comment 11, page 15)</p>	<p>Reconsider the change of categorisation of computing services (“capability and capacity computing” vs. “HPC and commodity computing”) as it does not seem to be fully successful. The diversity of the e-Infrastructure component ecosystem should be ensured.</p>	<p>Response: The categorisation of computing solutions was indeed a challenging area, and complicated by developments in both technical and organisational domains. In the end, the separation of HPC services (especially high-end HPC) from solutions where commodisation has progressed further seemed to be the most pragmatic one – but the categorisation of computing services will undoubtedly evolve (perhaps also through new solutions anticipated in the data-intensive computing domain). In any case, the intention of the roadmap was in no way to imply that the challenges related to large-scale installations are “easy” to solve in the context of any of the e Infrastructure services.</p> <p>Regarding the paragraphs identified in sections 2.2 and 2.5: one way to look at it is to see addressing the former as a precondition to achieve the latter – and also to achieve a balance between commercial markets and dedicated research e Infrastructure services. These issues would no doubt be studied in more detail by the suggested Task Force (referred to in the answer to Kees Neggers’s comment above).</p>
<p>Christos Kanellopoulos (comment 11, page 15)</p>	<p>Define “dedicated e-Infrastructure service”; outline challenges related to handling a large number of processing units; correct spelling issues.</p>	<p>Response: in this context “dedicated” denotes services provided by e Infrastructure-specific service providers (i.e. no commercial users), either to a specific research community or as a generic service to all research e Infrastructure users.</p> <p>Finally, the final version will be proofread by a professional editor, which should clearly improve readability of the text. Spelling will be based on UK spelling conventions, rather than US ones.</p>

Comments in the original form

1. Philippe Gavillet

Organisation - CERN & CIEMAT

Position - Senior Physicist

Dear Leif, We have read with much interest the e-Infrastructure Roadmap 2009 and acknowledge the impressive amount of work invested in the document. We basically agree with the various recommendations of the e-IRG with the two following minor comments: - Being members of a non-European e-Infrastructure (EELA-2) we much appreciate all actions to support International Collaboration spanning several continents, as expressed in 3.9. - As researchers we are rather sensitive to the organization of the computing of scientific groups, as discussed in sec.2.5. By experience we know that the existence of local, well adapted and supported ICT resources is mandatory to cover all computing needs of researchers groups. The part of the computing which can be outsourced to global e-Infrastructures is recognized. It concerns essentially CPU and Data intensive tasks. This outsourcing should not be at the expense of weakening the necessary local resources. Moreover, in many occasions a downsizing of the e-Infrastructure to run locally, in standalone, could be much better adapted to medium size scientific analyses. Best wishes Ph.Gavillet & B.Marechal (EELA-2).

2. Carlos M. Romeo-Casabona

Organisation - Inter-University Chair in Law and the Human Genome, University of Deusto and University of the Basque Country, Bilbao, Spain

Position – Director

Updating made in this new version seems to me to be adequate. Maybe Recommendation 3.8 should be rethought in a more realistic view: it could be not so easy to extend the provision of resources for some user communities of new EU member States.

3. Prof Alex Reid

Organisation - AARNet and University of Western Australia

Position - Advisor, eResearch & Middleware

Overall, we commend the latest version of the e-Infrastructure Roadmap In particular, we fully endorse the point, which the Roadmap makes quite forcibly (eg in §2.2), that researchers need infrastructure which in many areas exceeds currently available commercial services; in Australia, we have made this point repeatedly to government (now fully accepted), that the research sector needs networks which have 3 orders of magnitude greater capacity than the Broadband networks typically provided commercially.

Whilst we agree that there comes a time when special-purpose infrastructures need to be replaced by commodity services (§2.4.3), it is quite clear that this stage has not yet been reached with networking, and is unlikely to for at least 10 years - and may never, given science's constantly growing appetite for high volumes of data (such as the LHC and then the SKA), as exemplified by the growing adoption of the Fourth Paradigm for Research (data-intensive science, §2.3.1).

We also endorse the need for network providers more and more to provide seamless access to fibre circuits "on demand", and to integrate these services with conventional IP services (§2.4.1). And to embrace the findings of the various international projects exploring Next generation Internet technologies (eg FIRE, Federica, GENI, ION). In relation to the observation that increasingly heterogeneous network communities may require

new governance or funding models (§2.4.1), we would like to bring to eIRG's attention that AARNet has had 20 years of experience (as an NREN) in charging for network traffic use.

We strongly advocate strategies along the following lines for ensuring equitable approaches to network financing: 1. for small ad-hoc network use by 3rd parties, traffic charging on a volume basis is appropriate; 2. for more sustained network use by 3rd parties, or where there is significant strategic benefit in encouraging research collaborations, volume-based traffic charges represent a significant disincentive to such collaborations; in these cases a subscription-based charge should be employed. We also bring to eIRG's attention that we have developed a very robust governance model, where AARNet is owned and operated (and funded) by its members (the universities and research agencies); we therefore encourage further dialogue with us when and if 3rd-party use of NREN networks becomes an issue, requiring reconsideration of governance and financing arrangements.

4. Dietmar Erwin

Organisation - Forschungszentrum Jülich
Position - PRACE Deputy Project Manager

[N.B. the subheadings have been added in order to make references to specific comments more accurate in the table in the beginning of this document]

General Comments on the “e-IRG road map 2009” document. Draft dated December 18, 2009. From the text as written it is not clear who the intended audience is. I assume, it targets policy makers rather than technical experts. In any case I believe the following improvements should be made:

4.1. e-Infrastructure definition

Give a short clear definition of “e-Infrastructures” (the introduction ‘e-IRG and its mission’ might be a good place). Assuming that all readers share the same view can easily lead to misunderstandings.

The definition of the e-Infrastructure should also – at some point – include a definition of the components that form it. As a minor detail, it would be useful to define the relationship between the singular form of the term “e-Infrastructure” and the number of distinct e-Infrastructure services available today.

4.2. Improve the readability of the text and the clarity of the message.

There are several places where minor changes will achieve this. (The most economical way would be to make updates directly to the word document, not in the pdf. – I would be prepared to make the updates, since I have them already marked in my paper copy)

4.3. The bullet list in the executive summary

This should correspond to the nine recommendations in chapter 3. HPC (3.4) is missing also 3.7. The last bullet is important and should be given a sub-chapter in 3.

4.4. Commodity computing

The topic of ‘Commodity Computing’ is not defined precisely and not really convincingly introduced. Many of the aspects presented here have been introduced in the past as ‘Utility computing’, ‘Service on Demand’ etc with technical underpinnings like Grid computing and now Cloud computing.

It should also be noted that the strong link between the resource owners, users and middleware components in the HEP domain precedes the Grid concept.

4.5. Roadmap as part of the e-IRG processes

The link between the title 'e-IRG Roadmap' and the 'Recommendations' in chapter 3 should be made clearer. It is not obvious that this is really a roadmap.

Clarifying the role of the e-IRG roadmap as key result of the e-IRG process described in the section "e-IRG and its mission" should be considered.

4.6. HPC

Specific comments on HPC: PRACE is mentioned in 2.3.3. Its full name 'Partnership for Advanced Computing in Europe' should be given there. 3.4 spells HPC differently. It might be worth mentioning that PRACE is already one of the recommended components.

4.7. Use of terms

Inconsistent upper case/ lower case spelling of terms. Misleading or incorrect terms at several instances (see also 2.)

4.8. Data

I find it hard to believe that the flood of published knowledge in 2020 might rival the flood of raw data today.

5. Thomas Eickermann

Organisation - Forschungszentrum Juelich

Position - Head of networking division, PRACE project manager

I have just two comments (in addition to the extensive comments from Dietmar Erwin to which I had access):
- in 2.2 it is mentioned that other fields apart from networking have not been seen as supplying a services since recently. I think that HPC this is only true on a European scale. On a national scale, HPC has been provided as a service in many member states also since 20 years or more. - 2.3.2: it is referred to the "above mentioned Exascale project", but this project is mentioned here for the first time (and further in 2.3.3).

6. Andrew Lyall

ELIXIR Project Manager

The ESFRI Biology Research Infrastructures are of crucial importance to Europe because they play so strongly to the European Grand Challenges of 1) healthcare for an aging population, 2) a sustainable food supply, 3) a competitive pharmaceutical and biotechnology industrial sector and 4) protection of the environment. They will all need a data-infrastructure to ensure that they can interoperate, share data and avoid duplication and waste. This in turn is likely to be reasonably demanding of e-Infrastructure and it is important that these requirements are taken into account.

Modern biology is an information science and has evolved in very close association with the WWW. This is likely to continue and the need to move large datasets around Europe is going to increase and therefore make increasing demands of the network infrastructure.

Storage is also going to be important as biology is now producing dataset of the same order of size as physics.

HPC will be useful for a small percentage of biology computing, but most of it is data processing for which large blade farms are required (cf CERN).

These requirements are so important and the relevant communities do not have a history of working well together so I think a permanent body needs to be established to ensure that they meet regularly and come to understand each other better than currently.

If the next generation of e-Infrastructure is not designed with biology & medicine in mind then Europe will miss a great opportunity and risks falling behind the US and Japan as well as the emerging economies of India and China.

Andrew Lyall
ELIXIR Project Manager

7. Peter Wittenburg

[N.B. the subheadings have been added in order to make references to specific comments more accurate in the table in the beginning of this document]

I found the roadmap paper a good document to read in general and there are lots of very good statements and recommendations. Let me make some comments nevertheless.

7.1. Trust

Much is said about getting new user communities on board which is correct, but I miss one important word in this respect which is "trust". From all the many discussions we had in the MPG (also with industry!) the word "trust" and its many facets were in the focus. It has a technological facet which has to do with "reliability, availability, robustness", it has a real "trust" connotation in so far as for example "data is the fuel for the data driven science - what is required that researchers will indeed store their data at some center and carry out their work there". There are even more facets. A roadmap needs to refer to these aspects and say something about the sensitive and mechanisms that are required to achieve trust.

7.2. Service oriented architectures and commodity computing

It certainly is true that HPC is a very important aspect. Yet the number of researchers doing real HPC is a small group compared to the rest. What are the right environments for all the others to carry out their IT-based work. The paper makes proper statements about some aspects, but I did not find the notion of service-oriented architecture which are complementary to current grid solutions (at least the grid solutions so far do not help us a lot) and certainly complementary to HPC. In our domain we clearly foresee a distributed landscape of many "small" services for which we yet do not have a proper solution. If we refer to "grid" currently we are bound too much to solutions coming from domains that do not understand our domain. You use the term "commodity computing" to refer to this area in an abstract way, but it does not make this aspect sufficiently clear.

7.3. Mind as a grand challenge

The report speaks about the grand challenges at a certain place. I always ask myself why we ignore the mind and the societies. It is as if we do not notice that maintaining stability of life will also depend on whether we will manage to keep "minds" and "society" stable. The influences of innovation and globalization on minds and societies are enormous and we see that many people cannot follow. Some of us use the term "social climate change" which may indicate this fundamental problem.

7.4. 2-tier structure

I would like to see a note about what I may call a 2-tier structure. We have, stimulated by the ESFRI process, many RI initiatives working hard on understanding the needs of and solutions for their community. This is important since they understand the domain details and also build up a generation of experts. On the other side we have the common infrastructures. Yet we don't know where exactly the borderlines between these two are and this will change over time. We need space for a close interaction between these 2 tiers in concrete projects as we do it for example on a small scale: a) with DEISA a data replication project, b) with grid people a project to include web applications and web services in an AAI scheme, c) with GEANT/eduGain a project to come to a harmonized 7.614AAI which we can then use. To me these projects show how important this interaction is. If we do not have such projects at a sufficiently large scale, the common infrastructures take the risk that they do not offer the right interfaces/services/etc. Somewhere in the text you say correctly that the other way - communities work for themselves - also has many risks.

7.5. Commercial solutions

Commercial solutions will come when there is a market. Yet all we do has too much dynamics and uncertainties. Even for the area of repository systems we see that more solutions are being worked out, which obviously has to do with the fact that the current solutions are too limited to do real research. Just take the metadata issue: some system support Dublin Core, but this is from the world of librarians. The researchers need their own vocabulary and need flexibility. So perhaps we need to make a statement again about sensitivity to the research process and needs. It will take a little bit more time until smart people in industry come up with proper solutions. Also in this respect the word "trust" is relevant. Do researchers trust a commercial offer from a US company to store their data with them, if a government can always decide to get full access to all of it? The MPG researchers gave a clear answer "NO".

7.6. Software crisis

You mention the "software crisis" which is a very good point. Your notion however should also include the huge amount of access and processing software which is currently created for which we do not have any maintenance model. Much money which is currently invested will produce smart and trained young people which is good, but in terms of software I see a huge cemetery where smart code will just disappear.

7.7. Standards

On page 14 you make a statement about standards which to me seems a little odd. "Once interoperability has been reached, standardisation will have ...". Well it is standardization which will help us to achieve interoperability. Perhaps you see what I mean.

Hope this is of help and as said: in principle I liked to read the paper.

best
Peter

8. Mirco Mazzucato:

“It should be noted that the early-adopters’ – who were often also resource owners – efforts have driven the development of useful innovative open products, especially in the fields of the middleware, hardware choices and operating system support in the Grid projects. However, the “computing as a service” model has taken sometime to establish itself, but is accelerating as can be seen from the EGI Blueprint”

9. Edward Aronovich

Aiming the research infrastructure not only as network, computers and storage facilities but more likely as tools for research, there is need to develop some more general concept of service based infrastructure for research.

Adoption of Platform as a Service (PaaS) and Application as a Service (AaaS) models in addition to Infrastructure as a Service (IaaS), should be strongly stimulated and supported with the aim to increase the sustainability of e-Infrastructures and to identify and provide innovative solutions which could find a larger use in the society.

In this case, there is a need to prepare common interfaces for developers and tailored services for researchers that need applications or services. Commercial companies should be encouraged to offer competitive services for the research community based on competitive pricing. This will increase the interest in this infrastructure and will allow researchers to use optimized infrastructure.

10. Kees Neggers

After some further reflection, I would like to propose that in order to make the process of across the board commoditisation of computing services beneficial for all scientific users, e-IRG would recognize at the earliest opportunity the need for standards establishing interoperability and for organisational and financial models creating transparency.

11. Christos Kanellopoulos

Name: Christos Kanellopoulos

Organization: Aristotle University of Thessaloniki

Position: Technical Coordinator, Grid Operations Center - EGEE III NA5 Activity Manager

One of the axis of the 2007 Roadmap was the distinction between capability and capacity computing infrastructures. In the 2009 roadmap this terminology seem to have been dropped in favor of "HPC" and "Commodity Computing". Although the authors try to present the notion of "Commodity Computing", I believe in the end it is a bit confusing and could create the impression to the reader that one of the differentiating factors is the quality.

Another general note is that this Roadmap does not stress enough the need for interoperability within and across the various e-Infrastructures. The research communities need to be able to utilize the infrastructure as whole. For certain research activities "Commodity computing" might be the sensible answer, for others there might be a need for HPC resources. It is important that these different types of resources are presented in a

homogeneous manner to the researcher and this is something that will be mainly driven by the need to access data that might reside in different places or perhaps link different type of resources in a complex workflow. This is somehow addressed in the last paragraph of section 2.4.

In paragraph 2.2, it states "However attempting to maintain solutions that directly compete with commercial ones may waste considerable amounts of resources, which may be equally harmful when taking the long-term ability of the e-Infrastructure to fulfil advanced requirements of the user communities into account." Another point of view might be that having a rich ecosystem of competing solutions (both at the technological and the financial level) will lead into better solutions for the research communities. Of course there should be a balance in this. I believe the wording in the first sentence of the second paragraph in section 2.5 is more proper. Also "fulfil" should be fulfill".

In the same paragraph it goes on: "Current financing models are not sufficiently transparent to allow comparison of commodity services with the dedicated e-Infrastructure ones, which often complicates decision-making based on costs and benefits". The word dedicated, does it have the meaning of dedicated to a specific research community or something else?

In Section 2.3 it states "Harnessing a very large number of processing units efficiently in supercomputing applications is technically challenging". This challenge is not limited only to supercomputing applications. With the aggressive development of technologies such as GPGPUs, it will be common in the very near future to have thousands of processing units available to any application.