

Workshop opérations 2011

Journée "offre de services"

20/10/2011

Présents

- Jean-Michel Barbet (SUBATECH)
- Cécile Barbier (LAPP)
- Catherine Biscarat (LPSC)
- David Bouvet (CCIN2P3)
- Carlos Carranza (CPPM)
- Jean-Claude Chevalleyre (LPC)
- Frédérique Chollet (LAPP)
- Nicolas Clémentin (LUPM)
- Hélène Cordier (CCIN2P3)
- Stéphane Delmotte (LBBE/PRABI)
- Christophe Diarra (IPNO)
- Eric Fede (LAPP)
- Clément Gauthey (IBCP)
- Pierre Gay (Université Bordeaux)
- Christine Gondrand (LPSC)
- Edith Knoops (CPPM)
- Liliana Martin (LPNHE)
- Gilles Mathieu (CCIN2P3)
- Emmanuel Medernach (LPC)
- Geneviève Moguilny (IPGP)
- Jérôme Pansanel (IPHC)
- Guillaume Philippon (LAL)
- Denis Pugnère (IPNL)
- Marcin Radecki (PL-GRID/CYFRONET)
- Geneviève Romier (IDGC)
- Albert Shih (OBSPM)
- David Wallom (UK NGS)
- David Weissenbach (INSU01 – IPGP)

Agenda : <http://indico.in2p3.fr/conferenceTimeTable.py?confId=5551#20111020.detailed>

Introduction

Hélène Cordier, IN2P3-CC

<http://indico.in2p3.fr/contributionDisplay.py?contribId=1&confId=5551>

Hélène presents the agenda of the day, objectives of the workshop and expected outcome before presenting an overview of France Grid Operations over the past 12 months.

The UK NGS and national e-infrastructure

David Wallom, OeRC, University of Oxford, UK

<http://indico.in2p3.fr/contributionDisplay.py?contribId=2&confId=5551>

The NGS model should be seen through the perspective of a funding model which is quite different from the French model: all research programmes have the exact amount of funding they need to provide their science, including the funding for e-infrastructure. This comes to a project involvement model where communities come with funding for their research. Additionally, a lot of NGS partners are universities, which in the UK are not state bodies (as they are in France). There are now 28 partner institutions (there were only 4 in 2006). There is inherently then a strong links between providers and users. There are also strong connections between NGS and institutes/projects in the United States and in Asia Pacific.

The UK NGI is composed of NGS, GridPP and national parts of international projects. Initially, NGS was dedicated to non High Throughput Computing, which was dealt with by Grid-PP. All NGS users are within the national VO ngs.ac.uk – not in any EGI VOs...

The NGS VO can be seen as a national VO which acts as a catch all to give a “taste”.

There are some Site VOs, equivalent of our local/regional VOs, and project VOs, similar to EGI VOs model. There is a small number of accounts available, otherwise users should go through home institutions. NGS has no ownership of resources and operates clouds as a service out of institutes input. The sharing and determination of ownership of services is key. Central Services are run on 7.5 distributed FTE on virtual resources that are also distributed.

Several types of resources like HPC are manageable within the production infrastructure. NGS has developed strong links to ELIXIR communities.

Several identity managers are used such as SARoNGS/ use of Shibolet, Globus On Line.

NGS does differentiate central services, which it provides, and community shared services. “Cloud is a community shared service, just like computing, WMS, portals...”

Community shared services are decided by users: “We will not/cannot dictate what we make available. We have to let the communities bring their services to us, but use the same standards/interfaces that we already have.”

With regards to cloud, 4 members of staff are involved into the NGS Agile Deployment Environments. It is to be noted that “teaching” is one of the biggest cloud application, which was not necessarily foreseen. NGS have a fully implemented business model for cloud services.

About the quality of the service to provide and from the eucalyptus experience, DW mentions that “If you want to run a production cloud, it’s worth paying for the m/w you use”

Questions/Discussion

E.Fede: If I am a project with a part of my budget dedicated to support of application, do I have to do it myself or can I delegate to the NGS?

DW: We work with large institutional projects so if you are a small project the first point of contact should be them. We aim to support the aggregation of users into self organized communities. That’s where application support has to start first.

F.Chollet: Thinking of the way you are working with Grid-PP, it is not clear for me how this is realized. Has there been a collaborating agreement with Grid-PP?

DW: The NGS is running services for GridPP and vice versa. As in any project collaboration with the NGS, the aim is to have an MoU between the project and NGS as a central services provider. We are drawing this up at the moment.

F.Chollet: How do you deal with resource allocation? Ex of GridPP: allocation through pledges...

DW: within a research project, resource allocation processes are not our business. This is the responsibility of the community to define their process.

FC: so if I am a new community with no resources yet, what should I do?

DW: we have a small “taster” amount so that you can use it as a start. Then you should go to your institution, and that’s the institution that should negotiate resources

JM.Barbet: How is the impact of scientific papers measured?

DW: it is not really “measured”. We are working with communities that are generating papers, through a principle of acknowledgement

F.Chollet: What resources do you have to provide central services, and are these services centralized or distributed?

DW: 7.5 FTEs for running central services. These are NGS resources. Services are distributed between RAL, Manchester, Oxford, Daresbury.

H.Cordier: You don’t own any of the services, right?

DW: we (as NGS) do not really own any of the user facing services. All done through collaborations

H.Cordier: How do you manage to rely on self-sufficient communities?

DW: That is the consequence of a national policy and funding model. Scientific communities should be autonomous to do their research.

Services and Operations in Polish NGI

Marcin Radecki, ACC Cyfronet AGH, Poland

<http://indico.in2p3.fr/contributionDisplay.py?contribId=3&confId=5551>

PL-Grid is a national project which goal is to support national research. International collaboration is allowed if there is collaboration with Polish researchers. 3 of the 8 sites are connected to PL-GRID and not on EGI-Inspire. Job submission is followed through an SLA process which is tuned to users’ needs and implementation on several middlewares.

The resource allocation model involves that each site has autonomy in deciding resource allocation.

With regards to access to services, getting credentials has to be easy: the credentials are dedicated to PL-GRID and provide local account and access to a given mw’s resources, NB: grid users grow out of local users. A user portal is a must.

PL-Grid provides an online CA to use within PL-grid services.

SLA management: there are 2 types of SLAs implemented, with a notion of resource threshold.

To monitor SLAs, we need SLA awareness in the m/w

Scientific Application Monitoring: additional tests at job submission – check that all applications are available. This is possible because all sites support applications in the same way.

User support in PL-Grid includes application support and deployment of services requested by users.

It is organised into 3 lines of support. 1 out of those 3 FTEs can be dedicated to application porting. HPC and cloud interconnection with SLA Management is in progress.

Questions/Discussion

H.Cordier: What manpower is available for application porting and user-support?

MR: 3 FTE, amongst which 1 that can be dedicated to application porting if needed

D.Wallom: RA process seems fine with 5 sites. How do you see that if the infrastructure extends?

MR: that is a sensitive point. That'd indeed easy to do with sites funded by PL-grid, but difficult for other sites. However, the model gives them the ability to provide SLAs with their own users. We can only encourage sites to join this model

E.Fede: do you plan to test this SLA method beyond grids, i.e. on other infrastructures like cloud?

MR: some parts are generic enough to be applied to clouds, but not sure for some others. In Poland we are just starting with cloud experience, so this needs to be addressed later.

G.Mathieu: Sites are autonomous in RA: Is there any national policy for allocating grid resources?

MR: no clear policy from the Polish government. In Cyfronet we produce a yearly report of resource usage that integrates the project and discipline dimension. I never heard about any complaints on the current distribution.

We were thinking about some PL-Grid policy, in the form of recommendations to site. This will be based on the usage observed at the site. E.g. if a site has free resources they should accept more SLAs. Once we've reached a point where all sites are fully used then we can think of resource sharing.

D.Wallom: I understood you didn't use VOMS, is that right?

MR: we do use VOMS. When PL-Grid users want to register they use the user portal which is interfaced to VOMS. The tricky part with the user portal is that we support local accounts. So if a user requests an account then a local account is created on all institutes. If you create an account in PL-Grid and use gLite, you will end up on your account.

D.Weissenbach: What happens if some user that under unpredicted circumstances cannot execute his work during the timeframe allocated, and wants to complete it later?

MR: we have around 20 grants right now so we didn't have the case of users that couldn't complete their work during the allocated period.

What is foreseen is the other way round

D.Wallom: What is your sustainability long term with regards to staff, services and international participation?

MR: we are getting money from the ministry and the funding for the next years has been granted.

G.Mathieu: SLAs: what about communities that can't estimate their real needs? Is there a basic SLA?

MR: Currently SLAs are not mandatory, but we are moving towards a model where we will not provide any resources without SLA. This means there will be a very basic SLA for users that don't have specific needs.

Offre de services existante dans France Grilles

Hélène Cordier, IN2P3-CC

<http://indico.in2p3.fr/contributionDisplay.py?contribId=24&confId=5551>

National VO: How to get beyond the initial set of services is tricky. There have been many discussions and the answers are not yet conclusive.

Questions/Discussion

D.Weissenbach: There is some additional service provided around metadata.

D.Wallom: FG probably has to define which products and services it wants to be responsible for. You don't have unlimited amount of manpower so you can't do everything. You can then approach communities with recommendations. Of course, communities will want as much as possible so it is important to make clear what you are ready to provide.

→ Necessity of a catalogue of services

D.Wallom: we have some regional grids which mainly follow the sets of NGS services

JM.Barbet: difficulties understanding this catalogue. This should be user oriented and we see a lot of things directed to sites. So are we clients or are we providers?

The sites should be hidden from the user view.

G.Mathieu: Sites are both users and providers, the scope of the different services has to be defined according to that.

D.Wallom: the ESFRI projects will be contacting all the NGIs to ask for support. So this needs to be clear. The job of the NGI is to support people to support others

G.Mathieu: We should not build 1 catalogue, but 2: one for user services and one for infrastructure services. This come down to the iceberg in the introductory slides: the visible and the hidden. We are not in a funding model where projects come along and pay for services, so we have to make this visible somehow.

D.Wallom: it is important to stress on the services provided to those institutions that provide user facing services, as your newer communities have resources that need to be "gridified".

Offre de service des grilles régionales

JC.Chevaleyre, Auvergrid - N.Clémentin, MSFG - J.Pansanel, IPHC - P.Calvat, TIDRA - P.Gay, MCIA-G.Philippon, GRIF

P.Gay: users are within the region, from all institutes in the region. We provide a grid site and a VO. Requirements collection is done on a case by case basis.

We are trying to attract users to us so we provide as much as we can, including porting effort.

We have only a very small number of users. They are new users, so they don't know what they need.

D.Wallom: UK big sites are in some ways similar to French regional grids: they're talking to new communities and get them on board. Yet, the real question is: what part of provided services could you provide at a national level?

In your regional grid you may not have access to specific resources. Being regional doesn't mean you won't specialize in the longer term.

P.Gay: I would expect from FG and/or a national VO is to be an ability to scale for users with a more mature project.

N.Clémentin: similar case in MSFG. We have to support new users, and national support could be nice because there are more people that can provide this support. Local msfg VO is difficult to manage because of that.

JC.Chevaleyre: Auvergrid is probably the oldest regional grid, and today its status is not clear. Since 1 year, the number of sites in Auvergrid has decreased and it is likely that the regional grid will get down to 2 sites in the near future.

G. Philippon: GRIF is more an LCG federation than a “real” regional grid, so the vision is different. Local VOs are not the most active part. We are mainly supporting existing VOs, and not really into the business of getting new users on board.

JP.Meyer: local VOs are actually “local sites” VOs, not GRIF-wide ones.

D.Wallom: It seems that those regional grids are operating at the “community services” layer while the role of the NGI should be at the “core services” layer.

In the UK we have regional grids but also some institutes that have campus grids. So there are a lot of similarities with regional grids here. We do differentiate what is a core service and what is a specific one – although training, outreach and support are cross-boundaries.

G.Romier: what are regional grids’ criteria for accepting global VOs and “established communities”?

G.Philippon: we accept all VOs but with basic support (best effort).

D.Wallom: our approach is very much on a site by site basis.

E.Fede: involved in the Rhone-Alpes VO. Part of the core services provided by IN2P3-CC, other part by IN2P3-LAPP. We do not provide these services for regional use, but for an international (ESFRI) project: CTA. From a site point of view the regional VO is not the main consumer of these services.

D.Wallom: this is the longer term model. The site supports services for a community it has interest in.

P.Calvat: TIDRA is the regional grid in Rhone-Alpes that has 5 sites. Most users don’t really need the grid, so this is not what we propose systematically. We provide data services such as SRM/LFC tools but also iRODS which is very powerful. We do have a lot of user oriented tools as well

G.Romier: how do you gather users’ needs? Do you use questionnaires?

P.Calvat: No, direct contact is the best solution for us. Users come to us.

D.Wallom: Similar method sometimes used at campus level in the UK. We try and work with users without altering to much their working methods.

J.Pansanel: IPHC is multidisciplinary. We have many different communities, including some people from universities. Some of these communities have no resources for IT so we are providing them with resources. Also provide iRODS.

F.Chollet: To me the strategy is still not clear. Why would FG try and pilot regional grids?

H.Cordier: this has to be looked at through a mutualisation perspective.

D.Weissenbach: does supporting a regional VO make us part of this regional grid? My understanding of regional grid is that they have been funded to get regional funding, that’s all.

Wrap up: what services offer can/should France Grilles provide?

Important highlighted needs – what would we expect from the NGI:

- Scalability of resources
- Mutualisation of expertise across regions
- Strengthening collaborations
- Mutualisation of support
- Coherent support of nation-wide (e.g. ANR) or international (e.g. ESFRIs) projects

D.Wallom said “regional grids are operating at the community services layer while the role of the NGI should be at the core services layer”. Is it interesting to think of our services as split into central, shared and specific? In our current view sites are both clients and providers. That could be a way of clarifying this. So, the answer is obviously “yes”, we just now need to start.

UKI NGI presents itself as the composition of NGS+GridPP+national parts of international projects. Could we translate that to the French NGI? Are we the sum of FG central services + LCG France + the French participation to international projects?

Is SLA management a part of what we could consider as service provision? Should we start going in that direction?

NGS demos

After lunch, David Wallom demonstrated the use of some NGS services, mainly around authentication and authorization:

- SARoNGS - <https://cts.ngs.ac.uk/example.html>
- NGS Portal - <https://portal.ngs.ac.uk/JobProfiles.jsf>
- Certificate Wizard - <http://www.ngs.ac.uk/tools/certwizard>
- GSI-SSHTerm - <http://www.ngs.ac.uk/tools/gsisshterm>
- UI/WMS - <http://www.ngs.ac.uk/ui-wms>

Forum troubleshooting

use-cases qui relèvent de la formation admin de haut niveau

Best practices cf ticket GGUS (69666) sur la non- soumission de jobs :

Vérifier les /etc/passwords pour affiner le diagnostic concernant le user mapping (en cas de soumission avec un rôle donné ou un autre) à confirmer avec F.Schaer et N. Clémentin selon le résultat (pb connu sur glast). Solutions à long terme à clarifier/ classifier: pas de généralisation

Ticket GGUS (75240): proxy longue durée à spécifier dans le VOMS.

Le renouvellement du service myproxy est intégré dans WMS, mais cela ne marche pas. Ch Diarra mentionne que cela fonctionnait auparavant avec un by-pass : donner le proxy à l'utilisateur – délégation de proxy direct. DW mentionne un pb avec le WMS de Marseille et le my proxy à GRIF. La procédure de vérification des VOs de scope local ou « national » est à établir. EF mentionne que certains VO managers ne sont pas réactifs « risque potentiel ».

Comment survivre au crash du serveur Quattor .

Quel est le status de ce service « maître » avec le SVN : Remonter un service svn et restaurer la base. C'est une question de définition de best practice et une politique de construction/ de sauvegarde au niveau du site. Contacter la liste quattor.

Architecture détaillant les services fonctionnant entre eux

Cela relève de la documentation/training

Présentation du modèle d'opérations

Gilles Mathieu, IN2P3-CC

<http://indico.in2p3.fr/contributionDisplay.py?contribId=6&confId=5551>

GM présente les idées et concepts des modifications du modèle d'opérations de France Grilles, entamées au printemps, ainsi que leur utilité pratique.

Validation des entités du modèle

Hélène Cordier, IN2P3-CC

<http://indico.in2p3.fr/contributionDisplay.py?contribId=7&confId=5551>

F.Chollet : Les tâches de coordination sont-elles celles appelées « missions permanentes » ?

H.Cordier : Pas uniquement. C'est une classification transversale

F.Chollet : Y a-t-il d'autres tâches mutualisées que celles liées à la coordination ?

G.Mathieu : oui, par exemple « administration de la nagios box » qui est une tâche mutualisée et partagée mais qui est clairement de l'administration de service, pas de la coordination stratégique.

F.Chollet : A la dernière réunion du CTE, j'ai émis une crainte sur le fait que tous les groupes d'expertise n'étaient pas représentés. Il est vrai que nous n'étions pas nombreux.

H.Cordier : c'est la conséquence de la phase de transition dans laquelle on se trouve – Il n'est pas encore forcément clair qui doit venir ou non, et c'est le but de la validation de ce jour.

G.Mathieu : le travail de communication est à faire pour signaler les agendas, et inviter les experts en fonction de l'actualité.

F.Chollet : C'est une bonne chose que les experts soient toujours sur la liste et sont toujours informés.

H.Cordier : les personnes nommées comme « adjoints » seront ajoutés à la liste pour être tenus au courant.

Suggestions pour l'agenda des prochaines visios CTE : Le modèle opérationnel, les métriques et la vision stratégique des groupes d'expertise par rotation.

Suggestions pour l'agenda des visios opérations :

- Inclure des présentations générales dans les visios qui suivent le CTE (type IRODS, HINTS)
- Inclure le suivi des tickets ou des problèmes opérationnels qui traînent, donner l'opportunité aux sites de faire remonter des problèmes particuliers
- Demander à l'avance si il y a quelque chose à faire remonter de manière à prévoir un créneau dédié dans l'agenda.

Conclusion - discussion

F.Chollet : est-on certains de pouvoir tout gérer au niveau des tâches mutualisées ?

G.Mathieu : c'est le but de la formalisation en cours, notamment en terme de criticité des tâches pour assigner des priorités.

F.Chollet : ce serait intéressant d'avoir une vision d'ensemble sur tous les groupes de ce qui est important, ce qui est mutualisé et si les gens qui remplissent ces missions arrivent à faire face. Cette vision d'ensemble manque encore, on va peut-être trop loin dans les détails.

G.Mathieu : Il faudrait établir une cartographie des missions au-delà du découpage par groupe, pour obtenir la vue globale. Mais cela ne pourra se faire que sur la base des tâches détaillées, qu'il faudra épurer pour en extraire l'essentiel.

JM.Barbet : Si la stratégie de FG est d'aller chercher des utilisateurs dans les pays émergents etc, et que la politique du site est beaucoup plus restreinte, comment le site s'approprie la stratégie de FG ?

H.Cordier : la stratégie de base évoquée il y a quelques mois est que FG a la vocation de conserver et de faire évoluer la grille opérationnelle de production. Il n'est pas question d'aller à l'encontre de la stratégie des sites

F.Chollet : FG est légitime à émettre une stratégie mais elle se doit de la rendre explicite, et elle doit essayer de suivre ses objectifs en capitalisant les efforts des sites.

G.Mathieu : la stratégie nationale n'a pas vocation à être appliquée telle quelle par l'ensemble des sites, c'est une somme.

F.Chollet : la clé de voute de la clarification d'une stratégie nationale sera la mutualisation.

H.Cordier : La question de la place des institutions dans France Grilles devrait être une conclusion de l'analyse en cours, avec la criticité des tâches spécifiées.

La réunion se termine à 18h15.