

Achievements and perspectives of the biomed technical team

Franck MICHEL¹

*CNRS/UNS, I3S laboratory, MODALIS team
06903 Sophia Antipolis Cedex, France
E-mail: fmichel@i3s.unice.fr*

Tristan GLATARD

*Université de Lyon, CREATIS; CNRS UMR5220; Inserm U1044; INSA-Lyon
Université Lyon 1, 69621 Villeurbanne, France.
E-mail: tristan.glatard@creatis.insa-lyon.fr*

Johan MONTAGNAT

*CNRS/UNS, I3S laboratory, MODALIS team
06903 Sophia Antipolis Cedex, France
E-mail: johan@i3s.unice.fr*

Operating a large production-level distributed computing infrastructures today still requires a costly human intervention to reach acceptable quality of service. This cost may be acceptable for scientific communities with solid IT support, but it often remains a show-stopper for others. The Life-Science Grid Community (LSGC) gathers five Virtual Organizations (VOs) related to the Life-Science field. Among those, the biomed VO has set up a technical support team to be the technical interface between VO users and the NGI sites providing computing, storage and infrastructure services. The support team is a front-line for handling requests from sites and users; it monitors the resources, and enforces a set of pro-active measures to improve the service quality. It liaises with EGI-Inspire instances (UCB, UCST) to report needs, share experience, and learn from other communities' experience and best-practices. This abstract presents those goals in more details, the achievements of the past year, and highlights current actions and challenges. Overall, the handling of technical issues in an international VO still requires substantial manpower, to the expense of domain-specific activities. Exciting challenges remain to bring the full potential of grids to the end users.

*EGI Community Forum 2012 / EMI Second Technical Conference,
Munich, Germany
26-30 March, 2012*

¹ Speaker

1. Introduction

Scientists organised in Virtual Research Communities (VRCs) and Virtual Organizations (VO) make use of subsets of resources allocated by National Grid Initiatives (NGI) that participate to the European Grid Initiative (EGI). For those communities, operating large subsets of this infrastructure with reasonable quality of service requires substantial efforts to cope with runtime issues such as configuration flaws, hardware failures, load balancing, or storage unavailability. While early grid adopters such as High Energy Physics (HEP) are able to invest the effort needed to operate large subsets of the infrastructure while ensuring quality of service [1], this effort is out of reach for more fragmented communities. Consequently, some of them have set up strategies to operate VOs with a reduced human cost. This paper focuses on the example of the biomed VO initiative to meet those needs.

2. Description of the Work

The biomed VO uses resources allocated by 150 resource providers from several NGIs. The status of each resource depends on the perspective it is considered, and it frequently happens that although a resource is be considered up and running from an administrator's perspective, it is not be working properly for users of a specific VO. From this observation arises the need for a support team to "take place in the VO user's seat", in order to monitor - on their behalf - the resource availability and quality of service at the grid infrastructure level. The biomed support team has been set up with this purpose, to provide VO-specific but application-independent support activity. It currently consists of eight teams of volunteers from the most active user groups in France, Hungary, Spain, Italy and Viet-Nam. The support is organized in duty shifts to provide sustained support with a reasonably limited burden on contributors. A Nagios box, operated by resource provider GRIF, monitors storage and computing resources as well as critical services like the LFC and VOMS. The team on duty applies documented daily tasks and procedures, and submits GGUS "team" tickets when appropriate. It identifies and follows up on issues, discusses salient technical problems and investigates solutions. It also represents VO users to discuss technical requirements with EGI-InSPIRE instances.

During the last 2 years, efforts have been devoted to organize the shifts, train team members, coordinate teams take-over, document daily tasks, procedures and advice on how to deal with typical issues on a dedicated team wiki². Lots of existing tools, services and portals are available to assist the support team in its tasks. A task initiated more recently has been to assess how those tools could be used by the support team, and how reliable the data they provide is with regards to the VO, to avoid misinterpretations. This includes but is not be limited to GSTAT³, GOCDB⁴, EGI Accounting Portal⁵, SAM⁶ MyEGI, VO Admin Portal⁷, and

² <http://lsgc.org/en/Biomed-Shifts:Index>

³ <http://gstat.egi.eu/>

⁴ <https://goc.egi.eu/>

⁵ <http://accounting.egi.eu/egi.php>

⁶ <https://wiki.egi.eu/wiki/SAM>

⁷ <https://vodashboard.lip.pt/>

Operations Portal⁸. A continuous work is also being done to improve existing procedures, develop new tools, extend the scope of the monitored resources, and figure out new metrics that allow to assess resource quality of service for biomed users. For the time being, the latter regards mainly the SEs which are the most critical elements from the user's perspective.

3. Impact

Typical life-science applications involve computational workflows consuming/producing large data sets. While a single job failure may be acceptable, a storage resource failure may be a show-stopper, in particular when the middleware does not provide a file replication strategy. Therefore, besides critical services (LFC, VOMS), the initial effort of the support team has specifically focused on Storage Elements (SE). Then, the monitoring has progressively spanned to computing services (CE and WMS). To complement Nagios probes, an on-going development effort consists in defining new ways of identifying and anticipating potential problems. Again, storage resources are the primary focus with attempts to detect erroneous data published in the BDII, detect full SEs and come up with ways to handle them, identify heaviest users, "dark" data (aka. "zombie" files) i.e. files stored on SEs but not registered in any file catalogue, or "ghost" files that have catalogue entries with no corresponding physical replica, etc.).

From March 2010 to December 2011 the VO support team has handled 415 GGUS⁹ tickets, that is an average of 5.1 tickets per week. This figure is an indicator of the high operational cost for an international VO. Over the last year the biomed VO consumed 12 millions normalized CPU hours. Usage data history shows that the average ratio of waiting jobs over running jobs is 3.5, that may be the cause for long delays in the jobs queues; a survey has just started to assess CEs reliability by other ways than SAM probes, and figure out possible reasons are currently investigated. During the same period, the used storage space has raised from approximately 1.2 PB to 2 PB, out of 3.6PB.

So far, collected experience has helped identify features that would improve support quality and efficiency, but that evaluated tools do not address. Some features are addressed by the LSGC (redundant VOMS server), while some are being discussed with the User Community Support Team (redundant LFC server, VO Operations Dashboard). The remaining features have been gathered in the specification of the LSGC Dashboard, still to be developed, with features such as the management of VO users life-cycle workflow, support for robot certificates, and VO specific accounting metrics.

A reflection currently questions the impact of the support team daily activity on the way the VO resources perform, and how VO users perception may change as a result of this activity [2]. Various criteria can potentially show a positive or negative effect, like the evolution of resource consumption and reliability, the pace of submitted and received GGUS tickets, or the evolution of user satisfaction. So far however, the hindsight is not sufficient to attest of significant trends, and those criteria do not prove to be meaningful enough yet.

⁸ <http://operations-portal.egi.eu/>

⁹ <https://ggus.eu/pages/home.php>

4. Conclusions

The biomed VO is a significant user of the computing and storage resources that are made available to its users. Resources are monitored, their scope is dynamically and seamlessly extended. As a result, computing power and storage is delivered at a production level to 295 registered users consuming 12 millions CPU hours yearly.

However, current support efforts mainly deal with generic operational tasks that remain a heavy burden. Data management definitely concentrates the most important effort, although substantial efforts are needed on the computing resource side, to develop more accurate monitoring tools, assess the need for more resources or ways to use them more efficiently. The daily workload required from support team members on duty is high, and the coordination of a collection of independent international teams proves to be challenging. The volunteering model may need to be questioned as the community grows, unless options to reduce the overall support cost and improve its efficiency be implemented. [2] analyses in more details the causes for the high workload in the context of the biomed VO, and discusses possible options. Such an option may be to encourage the convergence of heterogeneous applications on a few common frameworks, for instance a community-dedicated instance of a generic framework like DIRAC [3] could be deployed and maintained by the VO support team, while the support effort would be mutualised with other VOs of the community.

Support concerns remain bound to the infrastructure, to the expense of VO-specific activities and collaborations within biomed and with other LSGC VOs. As a result, little is invested on the support of usual life sciences applications, on the sharing of community-specific data, knowledge, tools and experiments, while this is certainly a key potential of grids. This remains an exciting challenge to tackle in the next years.

5. References

- [1] J. Shiers. "Grid today, clouds on the horizon". *Computer Physics Communications*, 180:559-563, 2009.
- [2] F. Michel, J. Montagnat, T. Glatard. "Technical support for Life Sciences communities on a production grid infrastructure". In Proceedings of HealthGrid'2012, Amsterdam, Netherlands, May 2012.
- [3] A. Tsaregorodtsev, N. Brook, A. Casajus Ramo, Ph. Charpentier, J. Closier, et al. "DIRAC3 - The New Generation of the LHCb Grid Software". International Conference on Computing in High Energy and Nuclear Physics (CHEP'09), Journal of Physics: Conference Series 219(6), 2010.