

Using VAPOR to improve VO administration and operations

Monday, 19 May 2014 16:15 (15 minutes)

VAPOR, the Vo Administration and operations PORTal, intends to help small to medium-size grid user communities to perform common VO administration and operational tasks, at a reduced human cost. Such communities may typically have no or few dedicated IT support, have fragmented user groups, and may use resources either dedicated or in an opportunistic manner.

A large variety of existing tools and portals are available to assist VO managers and support teams. Although generic, some of them are often designed to meet the needs of somewhat specific contexts, and their operation requires a solid IT support. VAPOR is designed as a generic, experiment-independent tool. It complements existing tools such as SAM Nagios and the VO Operations Portal with novel services, to assist community managers and support teams in performing their daily activities. It does not require any deployment, since an existing instance of VAPOR is ready to support new VOs.

In this session, we shall describe the architecture of the portal, the status of the project (that ended in April 2014), and demonstrate its major features:

(i) resources monitoring (reports of resources supporting the VO, status indicators, statistical reports, white list of computing elements), and (ii) community data management (detect and handle SE filling-up, clean-up files of former users, deal with dark data and lost files).

Finally, we shall discuss future works and features to be implemented.

Wider impact and conclusions

Operating large subsets of the European Grid Infrastructure with reasonable QoS requires substantial efforts to cope with runtime issues. As this cost is often out of reach for small and fragmented communities, VAPOR proposes to help sustain communities by mutualizing the administrative and operational cost at the level of a VRC or beyond, and facilitate the outreach of new user communities by making it easier to start with the administration and operations of a VO.

At the time of writing, VAPOR is deployed and used by the biomed support team. Early feed back indicates that several other VOs are very interested in using it (Auger, CompChem, France Grille VO, VLEMED, WeNMR). During year 2014, we expect to open VAPOR to those VOs, as well as new emerging communities. Besides, we think that promoting the set-up of a common support team would help to mutualized the support effort, share experience, and overall improve the experience of those user communities.

URL(s) for further info

Project wiki: https://wiki.egi.eu/wiki/VT_VAPOR

Production instance: <https://operations-portal.egi.eu/vapor?vo=biomed>

Description of work

VAPOR is a 12-month mini-project funded by EGI-Inspire, that ended in April 2014. The first phase was dedicated to refine the scope of the project and define the functional specifications of the features to be developed. This work was achieved in collaboration with all partners, along with a review of existing reusable material and software.

Three main axes are identified:

1. Resource status indicators and statistical reports: VO support teams need to have a clear overview of the VO activity in terms of resource status and usage. VAPOR generates on-demand reports to: (i) visualise data and status information of resources supporting a VO; (ii) monitor the running and waiting jobs to help monitor peaks of activity and figure out bottlenecks; (iii) monitor computing elements status and time response; (iv) compute a white list of computing elements to feed job submission systems.
2. VO data management features aim at providing tools to: (i) clean-up files left behind by former users; (ii) check and fix issues of consistency between the storage elements and the file catalogue (dark data, lost files);

(iii) monitor storage elements to prevent them from feeling up.

3. User community management: VAPOR proposes to implement a Users Database which goals are to (i) manage and follow up on users registration life-cycle; (ii) track information about users “hidden behind” a robot certificate; (iii) track information about scientific publications to encourage users to acknowledge the usage of EGI resources.

At the time of writing (January 2014), items 1 and a part of item 2 are completed. We are confident that item 2 will be completed by the end of the project, however it is unlikely that item 3 will be covered due to lack of time.

VAPOR has been deployed to support the biomed VO, and is now used in production by the biomed VO support team. It is planned to extend the support of VAPOR to candidate VOs during the first quarter of 2014.

Primary author: MICHEL, Franck (CNRS)

Presenter: MICHEL, Franck (CNRS)

Session Classification: Tools for operating Research Infrastructures