



CMS Answers to the LCG/EGEE Operational assessment questionnaire for VOs

Name of the Virtual Organisation:

CMS

Name and contact details of the maintainer:

The contacts and maintainers for the CMS VO Boxes and CMS specific services are site specific and well known to the site personnel, as they have been negotiating the installation of each of the CMS specific services with the site managers in the first place and function as the general CMS site contact. The names and contacts for each of the CMS Tier-1 sites are available on request.

The name of the CMS Computing Coordinator, as of this writing, is Lothar A. T. Bauerdick, and can be contacted at Lothar.Bauerdick@NOSPAMcern.ch

1. **What are the hardware requirements for your VO Box?** (*Example – A node comparable to your site's fastest batch workers is sufficient. Twenty GigaBytes of possible network attached disk space is required.*)

Any modern dual CPU system would be sufficient. The system would preferably be better than a 2.4GHz Xeon system or equivalent Opteron. At least 1GB of RAM is required. 20GB of attached disk space is required

2. **What operating system is preferred or would be acceptable?** (*Example – We prefer RHEL3 derivatives such as CERN Scientific Linux 3, Scientific Linux 3, CentOS 3. RHEL4 would be acceptable*)

We prefer CERN SL3 or SL3. We expect to migrate to a flavour of SL4 during 2006. A derivative of RHEL3 would be acceptable.

3. **What certification has been performed of the VO services?** (*Examples – The software has been used within the EGEE pre-production service. The software has been used within the ARDA prototype testbed. The software has been test deployed at 3 Tier1 centres. The software is yet to be written or has had no testing.*)

The MonaLisa software has been installed on the OSG sites as a production service. It has also been installed at selected LCG-2.x sites for testing and integration.

PhEDEx has been installed on approximately 40 CMS sites for use in production transfers and Service Challenge Transfers.

PubDB has also been installed on approximately 40 CMS sites to publish the existence of data for production jobs and Service Challenge testing.

BOSS has been used extensively in official CMS production. The current version of BOSS has been deployed at CERN and FNAL and a number of desktop systems.



4. **Can the VO Box for your VO be co-located on the same physical hardware as another VO's VO Box?** (*Example – This is not possible since we require our services to bind to a particular port that cannot be altered or can not bind to a particular network interface..*)

The VO box can be co-located in the same physical hardware with other VOs. Two of the services (Phedex and BOSS) rely on mySQL instances. The traffic on these can be substantial, so it may not be optimal to share the physical hardware with other VO services that require access to mySQL.

5. **Can one VO box be shared across multiple sites to be used by the VO?** (*Example – A VO box cannot be shared across sites since it requires access to the NFS shared software area not typically available on the WAN.*)

The VO box should not be shared across multiple sites because the PhEDEx service requires close connections to local mass storage. PubDB could be shared across several sites, but since the service publishes local information controlled at a site level it is generally hosted locally.

6. **Should any VO services be monitored by the site? If so please provide technical details of how this can be achieved. In the case where incidents are triggered please provide details of the action required?** (*Examples – Please use a TCP Connect check on port 8123 to check service X is listening. If this test fails please reboot the node and subsequently telephone us immediately at any time of night or day.*)

All services can be centrally monitored, and the services will be monitored by the CMS site maintainers and contact personnel.

7. **How should the VO box be recovered after a complete loss? If more than a return of the vanilla VO box service is required please include a list of directories to be backed up and if any provision should be made for preserving and restoring consistent databases. Once the service is restored to previous state are any subsequent actions required by a site such as synchronisation with job outputs that completed during the downtime?** (*Examples – A reinstallation of a vanilla VO Box with no VO additions is acceptable. A recovery from a nightly backup is required. All areas under the home directory should be recovered after a backup. MySQL tables should be write locked during a backup..*)

A reinstallation of the VO box with the backed up services is required. The mySQL files for the database need to be restored.

8. **What is the time period in which a production service can be restored by the VO?** (*Production service will be restored by the VO the next working day.*)

The CMS site maintainers or somebody else from CMS will restore production service on the next working day.

9. **Are there any special requirements for the service which might influence the VO boxes initial deployment? For instance are there special network requirements of the VO box? Please provide details of how a VO service fits in to the overall workflow for the experiment.** (*Examples – Only command and control operations will take place so just the standard production network is required. All data access will be channelled through the VO box. Please see a particular document for a description of the VO service in the workflow.*)



There are no special requirements for these services. Production services are sufficient.

10. **What is the business impact for the failure of VO Box?** (*Examples – All jobs at the site will be lost. We will be unable to submit new work to this particular site. We will be unable to access any existing data at this site.*)

The failure of the VO box will cause the site to cease to be useful for new processing and transfer requests. Processing requests in the running state will likely succeed and active transfers will complete. New requests will fail. Real time monitoring of some services and utilization will be missing during the failure.

11. **How should the VO be contacted for service interventions?** (*Examples – Via a mailing list. Via an EGEE broadcast. Via status tags in the GlueSchema.*)

Service Interruptions should be sent to the CMS Operations mailing list and published in the information system within the GlueSchema.

12. **Are any steps required before or after the service outage?** (*Examples - Please drain all work from your farm belonging to VO X before the VO box is rebooted.*)

Nothing particular required, and the CMS site maintainers are responsible for eventually required steps.