

LHCb VO-box requirements

LHCb requires a dedicated VO-box to be set up in each LCG T1 site. The VO-box is a host dedicated to LHCb, which is running LHCb specific services. The purpose of the VO-box is twofold:

- Distribution of some of the critical LHCb VO services to allow management of their high load (load balancing) and high level of availability due to extra redundancy;
- Performing asynchronous operations for the LHCb jobs running on the LCG Grid to recover some of the failed operations.

VO-box Services

Currently the LHCb VO-boxes are running the following services :

- Configuration Service (CS) – is part of the LHCb Configuration System, which provides redundancy for serving Configuration data to all the components of the DIRAC Distributed Computing system. The CS running on a VO-box is a slave server which is automatically updating its copy of the configuration data from the Master CS running at CERN. The clients of the Configuration System running anywhere on the Grid can interrogate any instance of the CS (Master or Slave) in a round robin way to ensure 100% availability of this service.
- RequestManagement Service (RMS) – is part of the LHCb Failover Management system. The RMS is getting requests, usually from running jobs, to perform certain operations that jobs failed to execute. Examples are : job state updates, data upload operations, and others. Typically, the RMS instances at VO-boxes at T1 sites are just relay points to move the requests to the central RMS at CERN from where the operations are executed in an optimized way. Running several instances of the RMS service is necessary for the high service availability.

Other services can be implemented and added later to run on the VO-boxes.

The services running at VO-boxes are all implemented using the DIRAC project secure Framework. All the client-service interactions are authenticated using the standard GSI security infrastructure. The data exchange is performed over encrypted channels. The client requests are logged on the VO-box service side to help any incident resolution.

VO-box requirements

VO-boxes at Tier1 sites are required to be accessible from all the LCG sites. There is no special relation of any Grid site and a particular VO-box, they all can be used in an interchangeable way. Therefore, certain ports on the VO-box must be open in the site firewall rules to any inbound connections. Currently, ports number 9131 and 9143 are required to be open for CS and RMS services respectively.

The VO-boxes are provided by the T1 centers as hosts with installed standard Grid OS and with the standard Grid User Interface. This includes installation and up to date maintenance of the directories with CA related CRL lists.

The LHCb software installation and maintenance is under full responsibility of the LHCb managers. Therefore, LHCb user with certain certificates/roles are allowed to do remote login to the VO-box hosts to perform software maintenance.

The Services on the VO-boxes are running in the user space meaning that the service processes are executed by a non-root user. However, in order to perform GSI authentication of the client queries, the services must have access to the service certificate issued by the VO-box corresponding authority. Since there is no clear definition of the service certificate, the LHCb VO-boxes are using host certificates copied to the DIRAC Service installation area with read access permission for the Unix user running the services.

Even if the VO-box services are run by a non-root user, the first time service installation needs root rights in order to set up the automatic services start at boot time. This is achieved by putting a special command to the */etc/inittab* file.

The backup of the installed software or other files generated by the services is not required. This might change in the future with the introduction of new services.

The requirements for the VO-box hardware have nothing special. A reasonable mid-range modern server will be acceptable. For example, 1 or 2 core 2GHz CPU with 2-4 GB memory and 200GB hard disk. It should have also a reasonable connectivity to the Wide Area Network.