

Proposal for OSG Application Software Installation Service (OASIS)

June 19th, 2012

1 Motivation

The existing mechanism for installing software at a site is a writable directory that is shared with all worker nodes. Software for a given VO is installed by sending installation jobs to all sites supporting that VO. These installation jobs are sent under special credentials mapped to an UNIX account with permissions to write in the OSG_APP area.

Problems:

- a new site needs the software to be installed,
- the site needs to create and maintain the special UNIX account,
- no quota policies can be enforced,
- common tools are often installed in multiple places,
- unless sites provide for dedicated WNs, or specific batch policies, installation jobs may need to wait in a queue.

A new mechanism, allowing the distribution of the new software after a single installation, is desirable.

2 Architecture

The proposed architecture for OASIS is a server-client model. Installation is performed once, at a central place, and the software is distributed automatically to all sites.

The proposed architecture implies that software is installed in a single place (or a reduced number of places), and that the software will be automatically distributed to all sites. This dedicated area is otherwise immutable and will not accept any other type of input, since the installation procedure requires authentication and authorization in order to submit jobs that update or modify the software store.

The software and data deployed at the server is automatically distributed to client sites. Each site should run a replica service, and one or more Squid caches. Clients should trust content based on X509.

The list of server sites (including both multipurpose ones and VO specific) has to be published. The Grid Operation Center (GOC) would include this information as part of OIM.

The proposed framework for this architecture is CVMFS.

2.1 Login

The proposed mechanism allows login via a gatekeeper. Only credentials carrying a special VOMS attribute should be accepted. When a new VO is created, the authentication and authorization mechanism in place at that site should add the appropriate new accounts for that VO's authorized users. Each VO user's credentials should be mapped to a different UNIX account, so that their installation jobs do not interfere with those being performed by other VOs. This also implies that, for new VOs, the corresponding UNIX account must be created.

If a special software area for common tools is to be shared between multiple VOs (like ROOT or GEANT for HEP communities, or BLAST for Biocomputing communities), credentials with an special VOMS attribute (shared by the VOs) is needed, and these credentials are mapped to a dedicated special UNIX account.

The alternative would be that all UNIX accounts have permissions to write in this common area. But in that case, if a VO is using their own central service and therefore is not allowed to use their own area at GOC, then they could not interact with this common area either.

A dedicated area could be used to distribute the OSG WN middleware too.

The gatekeeper gives access to a local batch queue with a very few nodes, managed by condor. The installation includes a dedicated condor job wrapper script (pointed by condor configuration variable USER_JOB_WRAPPER). Documentation on this wrapper script can be found at

<http://research.cs.wisc.edu/condor/manual/v7.6/3.3Configuration.html#19134>

This wrapper script performs a set of extra tasks:

- writes a logbook recording information like: identity (DN), timestamp, change in service catalog, etc.,
- checks if there is enough space available for that VO to proceed,
- runs an atomic publishing command after installation process is complete,
- etc.

If needed, that wrapper can check the VO the job belongs to and call a dedicated vo-wrapper with specific tasks VO-specific.

3 Policies

3.1 Space quota

Each VO should be guaranteed the same amount of space when using the central service provided by OSG at GOC. If a VO needs additional space, it should consider the deployment of its own, separate repository service. In that case, that VO should not be allowed to have space at the central service (but would still be allowed to manage the common area).

3.2 Content

Only software with public licenses can be deployed. The VOs must agree to permit OSG the right to inspect the content of the software to ensure that their software and data are in compliance with relevant rules and laws.

4 Recommended procedures

In order to avoid inconsistencies between different replicas of the same file, it would be desirable to prevent the deletion and/or rewriting of files at the code source site.

If file deletion and re-writing are forbidden, all files must have a clear, well defined version number to distinguish one version from another.

5 OASIS distribution

OSG will publish the packages (RPMs, configuration files, etc.) for the clients and for the server for those VOs interested in running its own one.