



Contribution ID: 112

Type: **Oral Presentation**

Worker Node Software Management: the VO perspective

Tuesday, 12 April 2011 16:30 (30 minutes)

Overview

Grid jobs often run software that is domain-specific, VO specific and sometimes even user-specific. They depend on a software stack beyond what is generally by default available on grid nodes. During the Dutch VL-e project an additional distribution of software packages was created, the Proof of Concept (PoC), which is also distributed to the WNs on the BiG Grid infrastructure (Dutch NGI). However this was still very generic and it had very long update cycles and needed to be done by site administrators. For scientific experiments more specific stacks are required with shorter cycle times and with less (or no) participation of site administrators.

The toolset that we have developed offers centralized package management complete with dependency resolution, without the need for administrator's privileges. It uses the VO software area, which makes packages available on all the worker nodes.

Impact

Maintaining dependencies of software is a time-consuming task, but it is a very generic problem at the basis. By pooling from a large available set called pkgsrc[2] and building the tooling for installation and maintenance on grid sites, the problem is reduced to defining packaging instructions and dependency resolution for the domain-specific parts. Any VO that needs to maintain a local software stack can adopt this method.

By transferring management of software packages from the infrastructure providers to the VO managers, the responsibility for the maintenance of the stack now lies with those who have the proper domain knowledge.

The toolset that we have developed offers centralized package management complete with dependency resolution, without the need for administrator's privileges. It uses the VO software area, which makes packages available on all the worker nodes.

Additionally, these methods are not limited to VO managers. Individual grid users can make use of the same tools to fit their grid jobs with the right environment, although they will have to do this on a job-by-job basis as they lack the privilege to make such environments persist between jobs.

1. <http://www.netbsd.org/docs/software/packages.html>

Description of the work

The Grid is a (potentially) heterogeneous environment. Therefore simply distributing binaries is not always sufficient. We have therefore chosen to use a model where binaries are compiled remotely per site. This assumes that all resources in one site are homogeneous.

This very issue comes down to dependency management. Our solution allows a VO manager to maintain a stack of software libraries independent of what is provided by the WN software, so it can be updated asynchronously.

Also, multiple versions of libraries can be maintained side-by-side without interference through the use of the Environment Modules package.

With the help of the Environment Modules package[1] we manage all the paths and environment variables that are needed so that the software can be correctly found and executed by the job.

Further tooling is developed to update, monitor and test installations at multiple sites, so the users can check the validity of all systems in a single view and decide which sites can run their jobs. All these tools run as normal grid jobs that are submitted to the site with the permissions of the VO manager.

1. <http://modules.sourceforge.net/>

Conclusions

We have presented a solution that allows VO managers to maintain a stack of software applications and libraries independent of what is provided by the WN software.

Although we have a working prototype, there are still some improvements that we want to make to address the following issues.

Many VO's in the EGI project might have similar software wishes, but because of communication barriers those similarities might be unknown. We therefore have some thoughts about a more collaborative distribution model that enables reuse of porting effort.

Our current implementation is only tested within the infrastructure of one NGL. We are therefore looking forward testing our software with a VO that uses resources from multiple countries.

Primary authors: VAN DOK, Dennis (Nikhef); SANTCROOS, Mark (Academic Medical Center Amsterdam)

Co-authors: Prof. VAN KAMPEN, Antoine (Academic Medical Center Amsterdam); KEIJSER, Jan Just (Nikhef); Dr D. OLABARRIAGA, Silvia (Academic Medical Center Amsterdam)

Presenter: SANTCROOS, Mark (Academic Medical Center Amsterdam)

Session Classification: User Support Services

Track Classification: User Support Services - Infrastructure