

# Enabling Virtualization for Atlas Production Work through Pilot Jobs

Monday, 23 March 2009 08:00 (20 minutes)

Omer Khalid, Paul Nillson, Kate Keahey, Markus Schulz

Given the proliferation of virtualization technology in every technological domain, we have been investigating on enabling virtualization in the LCG Grid to bring in virtualization benefits such as isolation, security and environment portability using virtual machines as job execution containers.

There are many different ways to go around about it but as our workload candidate is Atlas experiment, so we choose to enable virtualization through pilot jobs which in Atlas case is Panda Pilot Framework. In our approach, once a pilot would have acquired a resource slot on the grid; it verifies if the server support virtual machines. If it does, then it proceeds to standard phases of job download and environment preparation and finally deploy virtual machine.

We have taken a holistic approach in our implementation where all the I/O takes places outside of the virtual machine on the host OS. Once all the data have been downloaded, then the Panda Pilot packages the job in the virtual machines and launches it for execution. Upon termination, panda pilot running on the host machine updates the server and stores the job output to an external SE and then do the clean up to makes the host slot available for next job execution.

Installing and maintaining Atlas releases on the worker nodes are the biggest issue, and especially how they could be made available to the virtual machine job execution container. In our implementation, Panda pilot takes an existing Atlas release installation and packages it in the virtual machine before starting it as read-only block device thus enabling the job to execute. Similarly, the base images for the virtual machine are generic to make sure that they are usable for large sets of jobs while keeping the control in the hands of system administrators as Panda pilot only uses the images made available by them.

In this way, pilot never loses the slot but at the same time enables virtualization on the grid in a systematic and coherent manner. Additional advantage of this approach is that only the computational over head of the virtualization is incurred which are minimal, and avoids more significant over head of I/O in a virtual machine by downloading/uploading in the host environment rather than in the virtual machine.

## Presentation type (oral | poster)

oral

**Primary author:** Mr KHALID, omer (CERN)

**Presenter:** Mr KHALID, omer (CERN)

**Session Classification:** Poster session

**Track Classification:** Distributed Processing and Analysis