

# Running jobs in the Vacuum

Monday, 16 September 2013 09:00 (8h 30m)

## Description of Work

An implementation of the vacuum scheme, Vac, is presented in which a VM factory runs on each physical worker node to create and contextualise its set of virtual machines. With this system, each node's VM factory can decide which VO's virtual machines to run, based on site-wide target shares and on a peer-to-peer protocol between factories. The site's VM factories query each other to discover which virtual machine types they are running, and therefore identify which virtual organisations' virtual machines should be started as VM slots become available again. This allows sites to provide virtual environments to VOs and still maintain fair share allocation of capacity between multiple VOs, as is currently achieved with grid interfaces to conventional batch systems. Another property of this system is that there is no gate keeper service, head node, or batch system accepting and then directing jobs to particular worker nodes, avoiding several central points of failure.

Finally, we describe use of the Vac system to run production jobs from the central LHCb task queue, using the same contextualisation procedure for virtual machines developed by LHCb for IaaS clouds and for BOINC, another IaaS system harnessing otherwise idle resources.

## Printable Summary

We present a model for the operation of computing nodes at a site using virtual machines, in which the virtual machines (VMs) are created and contextualised for virtual organisations (VOs) by the site itself. For the VO, these virtual machines appear to be produced spontaneously "in the vacuum" rather than in response to requests by the VO. This is an inversion of the usual cloud model of Infrastructure-as-a-Service, and these sites operate as Infrastructure-as-a-Client (IaaSC), with work pulled from VO clouds rather than sent to site clouds.

This model takes advantage of the mature pilot job frameworks adopted by many VOs, in which pilot jobs submitted via a grid infrastructure in turn start job agents which fetch the real jobs from the VO's central task queue. In the vacuum model, the contextualisation process starts a job agent within the virtual machine and real jobs are then fetched from the central task queue as normal for these systems. This parallels similar developments in which cloud interfaces are used to start virtual machines containing job agents.

**Primary author:** MCNAB, Andrew (MANCHESTER)

**Co-authors:** STAGNI, Federico (CERN); UBEDA GARCIA, Mario (CERN)

**Presenter:** MCNAB, Andrew (MANCHESTER)

**Session Classification:** Posters display