



Contribution ID: 113

Type: oral presentation

Pilots 2.0: DIRAC pilots for all the skies

Monday, 13 April 2015 14:00 (15 minutes)

In the last few years, new types of computing infrastructures, such as IAAS (Infrastructure as a Service) and IAAC (Infrastructure as a Client), gained popularity. New resource may come as part of pledged resources, while others are in the form of opportunistic ones. Most of these new infrastructures are based on virtualization techniques, others don't. Meanwhile, some concepts, such as distributed queues, lost appeal, while still supporting a vast amount of resources. Virtual Organizations are therefore facing heterogeneity of the available resources and the use of an Interware software like DIRAC to hide the diversity of underlying resources has become essential.

The DIRAC WMS is based on the concept of pilot jobs that was introduced back in 2004. A pilot is what creates the possibility to run jobs on a worker node. The advantages of the pilot job concept are now well established. The pilots are not only increasing the visible efficiency of the user jobs but also help to manage the heterogeneous computing resources presenting them to the central services in a uniform coherent way. Within DIRAC, we developed a new generation of pilot jobs, that we dubbed Pilots 2.0. Pilots 2.0 are not tied to a specific infrastructure; rather they are generic, fully configurable and extendible pilots. A Pilot 2.0 can be sent, as a script to be run, or it can be fetched from a remote location. A pilot 2.0 can run on every computing resource, e.g.: on CREAM Computing elements, on DIRAC Computing elements, on Virtual Machines as part of the contextualization script, or IAAC resources, provided that these machines are properly configured, hiding all the details of the WN's infrastructure. Pilots 2.0 can be generated server and client side. Pilots 2.0 are the "pilots to fly in all the skies", aiming at easy use of computing power, in whatever form it is presented. Another aim is the unification and simplification of the monitoring infrastructure for all kind of computing resources by using pilots as a network of distributed sensors coordinated by a central resource monitoring system.

Pilots 2.0 have been developed using the command pattern: each command is realizing an atomic function, and can be easily activated and de-activated based on the WN type. VOs using DIRAC can tune pilots 2.0 as they need, and extend or replace each and every pilot command in an easy way. In this paper we describe how Pilots 2.0 work with distributed and heterogeneous resources providing the abstraction necessary to deal with different kind of computing resources.

Primary author: STAGNI, Federico (CERN)

Co-authors: Dr TSAREGORODTSEV, Andrei (CPPM, Aix-Marseille Université, CNRS/IN2P3, Marseille, France); MC-NAB, Andrew (University of Manchester (GB)); LUZZI, Cinzia (CERN)

Presenter: STAGNI, Federico (CERN)

Session Classification: Track 4 Session

Track Classification: Track4: Middleware, software development and tools, experiment frameworks, tools for distributed computing