

hybrid batch system deployment with AWS spot instances

Monday, September 23, 2019 4:45 PM (15 minutes)

Our institution, Port d'Informació Científica (PIC), is an innovative centre for supporting research and provides support to scientific groups working in projects which require large amount of computing resources for the analysis of massive sets of distributed data. PIC is the Spanish Tier-1 center for the Large Hadron Collider, the main (Tier-0) data center for the MAGIC telescopes and the PAU dark energy survey, and is one of the Science Data Centers of ESA's Euclid mission.

At PIC we have piloted a hybrid cloud computing platform totally integrated in our batch computing service and transparent to the final users. We doubled our computing capacity using AWS spot instances for 72 hours in order to test how we can increase our peak computing needs at an affordable price.

To test this hybrid batch system infrastructure we have used the HTCondor condor_annex tool, which makes the process of extending a local pool with cloud resources easy, fast and if the user needs it, with an expiration date. In order to get to the production ready system, everything was tested in three steps: small batch of on-demand instances in a test environment, small batch of on-demand and spot-instances in a production environment and big batch of spot instances in a production environment.

Initially the jobs were sent to our test environment to then be moved to production after checking that the jobs were running correctly, both of them using on-demand instances. The test continued by launching spot-instances in a seamless hybrid infrastructure where the cloud worker nodes were added to the local computing pool and have jobs running in minutes. Accounting and monitoring of the cloud resources has been totally integrated with the local system.

Amazon Web Services Spot Instances offers the possibility to instantiate machines at a fraction of the on-demand price due to low demand of specific instance types at specific times. When a lot of instances are launched and the conditions to keep them running change, some or all of them can be stopped at any moment. This suits very well use cases like the one tested here.

There were some other elements needed to configure the system, such as a custom worker node image created and stored in a specific region in AWS or a HTCondor Connection Broker (CCB) to enable communication between the AWS nodes and the local system, apart from the changes in the HTCondor configuration to accept the new servers as own.

Primary authors: Dr ACOSTA, Carles (IFAE); Dr MERINO, Gonzalo (CIEMAT); Mr CASALS, Jordi (CIEMAT); Ms ACÍN, Vanessa (IFAE)

Presenter: Mr CASALS, Jordi (CIEMAT)

Session Classification: IBERGRID Contributions

Track Classification: R&D for computing services, networking, and data-driven science.