CHEP 2016 Conference, San Francisco, October 8-14, 2016

Contribution ID: **511** Type: **Poster**

Plancton: an opportunistic distributed computing project based on Docker containers

Thursday, 13 October 2016 16:30 (15 minutes)

The computing power of most modern commodity computers is far from being fully exploited by standard usage patterns.

The work we present describes the development and setup of a virtual computing cluster based on Docker containers used as worker nodes. The facility is based on Plancton[1]: a lightweight fire-and-forget background service that spawns and controls a local pool of Docker containers on a host with free resources by constantly monitoring its CPU utilisation. Plancton is designed to release the resources allocated opportunistically whenever another demanding task is run by the host user, according to configurable thresholds: this is attained by killing a number of running containers.

The resources comprising the facility are a collection of heterogeneous non-dedicated Linux hosts ideally inside the same local network, with no guaranteed network bandwidth, made available by members of a collaboration or institute. The user has agreed to donate its spare CPU cycles and remains the administrator of the involved host. Since the system is based on Docker containers performance isolation and security are guaranteed through sandboxing.

Using a thin virtualization layer such as Docker has the effect of having containers that are started almost instantly upon request. We will show how fast startup and disposal of containers finally enables us to implement the formation of the opportunistic cluster in a headless fashion, where our containers are mere pilots.

As an example we are running pilot HTCondor containers automatically joining a given cluster and terminating right after executing a job or in a short while if no new job is available. Software is provided through CVMFS on Parrot, making the execution environment suitable for HEP jobs.

Finally we will show how the uncomplicated approach of Plancton to containers deployment makes it suitable for setting up dedicated computing facilities too, provided that the underlying use case is sufficiently simple.

[1] https://github.com/mconcas/plancton

Tertiary Keyword (Optional)

Distributed workload management

Secondary Keyword (Optional)

Virtualization

Primary Keyword (Mandatory)

Computing middleware

Primary author: CONCAS, Matteo (Universita e INFN Torino (IT))

Co-authors: BERZANO, Dario (CERN); MASERA, Massimo (Universita e INFN Torino (IT)); PUCCIO, Maximiliano (Universita e INFN Torino (IT)); VALLERO, Sara (Universita e INFN Torino (IT)); BAGNASCO, Stefano (I.N.F.N. TORINO); LUSSO, Stefano (INFN-TO)

Presenter: CONCAS, Matteo (Universita e INFN Torino (IT))

Session Classification: Posters B / Break

Track Classification: Track 7: Middleware, Monitoring and Accounting