

Using container orchestration to improve service management at the RAL Tier 1

Monday 10 October 2016 15:45 (15 minutes)

At the RAL Tier-1 we have been deploying production services on both bare metal and a variety of virtualisation platforms for many years. Despite the significant simplification of configuration and deployment of services due to the use of a configuration management system, maintaining services still requires a lot of effort. Also, the current approach of running services on static machines results in a lack of fault tolerance, which lowers availability and increases the amount of manual interventions required. In the current climate more and more non-LHC communities are becoming important, with the potential for the need to run additional instances of existing services as well as new services, but at the same time comes the likelihood that staff effort is more likely to decrease rather than increase. It is therefore important that we are able to reduce the amount of effort required to maintain services whilst ideally improving availability, in addition to being able to maximise the utilisation of resources and become more adaptive to changing conditions.

These problems are not unique to RAL, and from looking at what is happening in the wider world it is clear that container orchestration has the possibility to provide a solution to many of these issues. Therefore last year we began investigating the migration of services to an Apache Mesos cluster running on bare metal. In this model the concept of individual machines is abstracted away and services are run on the cluster in Docker containers, managed by a scheduler. This means that any host or application failures, as well as procedures such as rolling starts or upgrades, can be handled automatically and no longer require any human intervention. Similarly, the number of instances of applications can be scaled automatically in response to changes in load. On top of this it also gives us the important benefit of being able to run a wide range of services on a single set of resources without involving virtualisation.

In this presentation we will describe the Mesos infrastructure that has been deployed at RAL, including how we deal with service discovery, the challenge of monitoring, logging and alerting in a dynamic environment and how it integrates with our existing traditional infrastructure. We will report on our experiences in migrating both stateless and stateful applications, the security issues surrounding running services in containers, and finally discuss some aspects of our internal process for making Mesos a platform for running production services.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Primary Keyword (Mandatory)

Cloud technologies

Primary author: LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB))

Co-author: COLLIER, Ian (STFC - Rutherford Appleton Lab. (GB))

Presenter: LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB))

Session Classification: Track 6: Infrastructures

Track Classification: Track 6: Infrastructures