

# Stealth Cloud: How not to waste CPU during grid to cloud transitions

*Thursday, October 13, 2016 12:00 PM (15 minutes)*

When first looking at converting a part of our site's grid infrastructure into a cloud based system in late 2013 we needed to ensure the continued accessibility of all of our resources during a potentially lengthy transition period.

Moving a limited number of nodes to the cloud proved ineffective as users expected a significant number of cloud resources to be available to justify the effort of converting their workflows onto the cloud.

Moving a substantial part of the cluster into the cloud carries an inherent risk, such as the cloud nodes sitting idle while waiting for the VOs to finish their development work and other external factors. To mitigate this, we implemented a system to seamlessly move some of the grid workload across to the cloud such that they could use any idle resources. A requirement for this was that the existing grid jobs must be transparently run in a VM without requiring any adjustments by the job owner. To accomplish this we brought together a number of existing tools, ARC-CE, glideinWMS (a pilot-based WMS developed by CMS) & OpenStack.

This talk will focus on the details of the implementation and show that this is a viable long-term solution to maintain resource usage during long periods of transition.

## Secondary Keyword (Optional)

Computing middleware

## Primary Keyword (Mandatory)

Cloud technologies

## Tertiary Keyword (Optional)

**Primary authors:** BAUER, Daniela (Imperial College Sci., Tech. & Med. (GB)); FAYER, Simon (Imperial College Sci., Tech. & Med. (GB))

**Presenter:** BAUER, Daniela (Imperial College Sci., Tech. & Med. (GB))

**Session Classification:** Track 6: Infrastructures

**Track Classification:** Track 6: Infrastructures