

Context-aware distributed cloud computing using CloudScheduler

Wednesday, 12 October 2016 12:45 (15 minutes)

The distributed cloud using the CloudScheduler VM provisioning service is one of the longest running systems for HEP workloads. It has run millions of jobs for ATLAS and Belle II over the past few years using private and commercial clouds around the world. Our goal is to scale the distributed cloud to the 10,000-core level, with the ability to run any type of application (low I/O, high I/O and high memory) on any cloud. To achieve this goal, we have been implementing changes that utilize context-aware computing designs that are currently employed in the mobile communication industry. Context-awareness makes use of real-time and archived data to respond to user or system requirements. In our distributed cloud, we have many opportunistic clouds with no local HEP services, software or storage repositories. A context-aware design significantly benefits the reliability and performance of our system by locating the nearest or optimal location of the required services. We describe how we are collecting and managing contextual information from our workload management systems, the clouds, the virtual machine and our services. This information is used not only to monitor the system but also to carry out automated corrective actions. We are incrementally adding new alerting and response services to our distributed cloud. This will enable us to scale the number of clouds and virtual machines. Further, a context-aware design will enable us to run analysis or high I/O application on opportunistic clouds. We envisage an open-source HTTP data federation (for example, the Dynafed system at CERN) as a service that would provide us access to existing storage elements used by the HEP experiments.

Primary Keyword (Mandatory)

Cloud technologies

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Primary author: Dr SOBIE, Randall (University of Victoria)

Co-authors: Mr LEAVETT-BROWN, Colin Roy (University of Victoria); Mr RING, Darryl (University of Victoria); Dr CASTEELS, Kevin (University of Victoria); Mr MURRAY, Matthew (University of Victoria); Mr PATERSON, Michael (University of Victoria); Dr SEUSTER, Rolf (University of Victoria); Mr TAYLOR, Ryan (University of Victoria)

Session Classification: Track 3: Distributed Computing

Track Classification: Track 3: Distributed Computing