20th International Conference on Computing in High Energy and Nuclear Physics (CHEP2013)



Contribution ID: 418

Type: Poster presentation

Dynamic web cache publishing for laaS clouds using Shoal

Monday 14 October 2013 15:00 (45 minutes)

It has been shown possible to run HEP workloads on remote IaaS cloud resources. Typically each running Virtual Machine (VM) makes use of the CERN VM Filesystem (CVMFS), a caching HTTP file system, to minimize the size of the VM images, and to simplify software installation. Each VM must be configured with a HTTP web cache, usually a Squid Cache, in proximity in order to function efficiently. Regular grid sites which use CVMFS to host worker node software areas use one or more Squid servers for their site. However, each IaaS cloud site has none of the static infrastructure associated with a typical HEP site. Each cloud VM must be configured to use a particular Squid server. We have developed a method and software application called Shoal for publishing squid caches which are dynamically created on IaaS clouds. The Shoal server provides a simple REST interface which allows clients to determine their closest Squid cache. Squid servers advertise their existence by running Shoal agent which uses the Advanced Message Queuing Protocol (AMQP)to publish information about the Squid server to the Shoal server. Having a method for exchanging the information rapidly allows for Squid servers to be instantiated in the cloud in response to load and for clients to quickly learn of their existence. In this work we describe the design of Shoal and evaluate its performance at scale in an operational system.

Authors: GABLE, Ian (University of Victoria (CA)); SOBIE, Randy (University of Victoria (CA))

Presenter: GABLE, Ian (University of Victoria (CA))Session Classification: Poster presentations

Track Classification: Distributed Processing and Data Handling A: Infrastructure, Sites, and Virtual-

ization