



Contribution ID: 227

Type: **Parallel**

The WorkQueue project - a task queue for the CMS workload management system

Tuesday 22 May 2012 13:30 (4h 45m)

We present the development and first experience of a new component (termed WorkQueue) in the CMS workload management system. This component provides a link between a global request system (Request Manager) and agents (WMAgents) which process requests at compute and storage resources (known as sites). These requests typically consist of creation or processing of a data sample (possibly terabytes in size).

Unlike the standard concept of a task queue, the WorkQueue does not contain fully resolved work units (known typically as jobs in HEP). This would require the WorkQueue to run computationally heavy algorithms that are better suited to run in the WMAgents. Instead the request specifies an algorithm that the WorkQueue uses to split the request into reasonable size chunks (known as elements). An advantage of performing lazy evaluation of an element is that expanding datasets can be accommodated by having job details resolved as late as possible.

The WorkQueue architecture consists of a global WorkQueue which obtains requests from the request system, expands them and forms an element ordering based on the request priority. Each WMAgent contains a local WorkQueue which buffers work close to the agent, this overcomes temporary unavailability of the global WorkQueue and reduces latency for an agent to begin processing. Elements are pulled from the global WorkQueue to the local WorkQueue and into the WMAgent based on the estimate of the amount of work within the element and the resources available to the agent.

WorkQueue is based on CouchDB, a document oriented no-sql database. WorkQueue uses the features of CouchDB (map/reduce views, bi-directional replication between distributed instances) to provide a scalable distributed system for managing large queues of work.

The project described here represents an improvement over the old approach to workload management in CMS which involved individual operators feeding requests into agents. This new approach allows for a system where individual WMAgents are transient and can be added or removed from the system as needed.

Student? Enter 'yes'. See <http://goo.gl/MVv53>

no

Summary

We present the development and first experience of a new component (termed WorkQueue) in the CMS workload management system. This component provides a link between a global request system (Request Manager) and agents (WMAgents) which process requests at compute and storage resources (known as sites). These requests typically consist of creation or processing of a data sample (possibly terabytes in size).

WorkQueue is based on CouchDB, a document oriented no-sql database. WorkQueue uses the features of CouchDB (map/reduce views, bi-directional replication between distributed instances) to provide a scalable distributed system for managing large queues of work.

The project described here represents an improvement over the old approach to workload management in CMS which involved individual operators feeding requests into agents. This new approach allows for a system

where individual WMAgents are transient and can be added or removed from the system as needed.

Authors: RYU, Seangchan (Fermilab); Dr WAKEFIELD, Stuart (Imperial College London)

Co-authors: Dr EVANS, Dave (Fermi National Accelerator Lab. (US)); Mr NORMAN, Matthew (University of California at San Diego); METSON, Simon (University of Bristol (GB)); FOULKES, Stephen (Fermi National Accelerator Lab. (Fermilab)); MAXA, Zdenek (California Institute of Technology (US))

Presenter: Dr WAKEFIELD, Stuart (Imperial College London)

Session Classification: Poster Session

Track Classification: Distributed Processing and Analysis on Grids and Clouds (track 3)