

Improved Cache Coherency Approach for CMS Frontier

Monday 23 March 2009 08:00 (20 minutes)

The CMS experiment requires worldwide access to conditions data by nearly a hundred thousand processing jobs daily. This is accomplished using a software subsystem called Frontier. This system translates database queries into http, looks up the results in a central database at CERN, and caches the results in an industry-standard http proxy/caching server called Squid. One of the most challenging aspects of any cache system is coherency, that is, ensuring that changes made to the underlying data get propagated out to all clients in a timely manner. Recently, the Frontier system was enhanced to drastically reduce the time for changes to be propagated everywhere, typically as low as 10 minutes for some kinds of data and no more than 60 minutes for the rest of the data, without overloading servers. This was done by taking advantage of an http and Squid feature called “If-Modified-Since” in which the “Last-Modified” timestamp of cached data is sent back to the central server. The server responds to this with a very short message if data has not been modified, which is the case most of the time, and re-validates the cache. In order to use this feature, the Frontier server has to send the “Last-Modified” timestamp, but that information is not normally stored by the Oracle databases so a PL/SQL program was developed to keep track of the modification times of database tables. We discuss the details of this caching scheme and the obstacles overcome including Oracle database and Squid bugs.

Authors: Dr DYKSTRA, David (Fermilab); Dr LUEKING, Lee (FERMILAB)

Presenter: Dr DYKSTRA, David (Fermilab)

Session Classification: Poster session

Track Classification: Distributed Processing and Analysis