Contribution ID: 494

Type: Oral

Use of DAGMan in CRAB3 to improve the splitting of CMS user jobs

Monday 10 October 2016 15:00 (15 minutes)

CRAB3 is a workload management tool used by more than 500 CMS physicists every month to analyze data acquired by the Compact Muon Solenoid (CMS) detector at the CERN Large Hadron Collider (LHC). CRAB3 allows users to analyze a large collection of input files (datasets), splitting the input into multiple Grid jobs depending on parameters provided by users.

The process of manually specifying exactly how a large project is divided into jobs is tedious and often results in sub-optimal splitting due to its dependence on the performance of the user code and the content of the input dataset. This introduces two types of problems; jobs that are too big will have excessive runtimes and will not distribute the work across all of the available nodes. However, splitting the project into a large number of very small jobs is also inefficient, as each job creates additional overhead which increases load on scheduling infrastructure resources.

In this work we present a new feature called "automatic splitting" which removes the need for users to manually specify job splitting parameters. We discuss how HTCondor DAGMan can be used to build dynamic Directed Acyclic Graphs (DAGs) on the fly to optimize the performance of large CMS analysis jobs on the Grid.

We use DAGMan to dynamically generate interconnected DAGs that estimate the time per event of the user code, then run a set of jobs of preconfigured runtime to analyze the dataset. If some jobs have terminated before completion, the unfinished portions are assembled into smaller jobs and resubmitted to the worker nodes.

Secondary Keyword (Optional)

Computing middleware

Primary Keyword (Mandatory)

Distributed workload management

Tertiary Keyword (Optional)

Authors: MASCHERONI, Marco (Fermi National Accelerator Lab. (US)); WOLF, Matthias (University of Notre Dame (US))

Co-authors: WOODARD, Anna Elizabeth (University of Notre Dame (US)); BOCKELMAN, Brian Paul (University of Nebraska (US)); VAANDERING, Eric (Fermi National Accelerator Lab. (US)); HERNANDEZ, Jose (CIEMAT); BELFORTE, Stefano (Universita e INFN, Trieste (IT))

Presenter: WOODARD, Anna Elizabeth (University of Notre Dame (US))

Session Classification: Track 3: Distributed Computing

Track Classification: Track 3: Distributed Computing