CENTRAL RECONSTRUCTION SYSTEM ON THE RHIC LINUX FARM IN BROOKHAVEN LABORATORY

A.Chan, R. Hogue, C. Hollowell, R. Popescu, G. Tsai, A.Withers, T.Wlodek *

Brookhaven National Laboratory, Upton, NY 11973, USA.

Abstract

This document describes the design and requirements of the Central Reconstruction System (CRS) of the Rhic Computing Facility (RCF) Linux farm in the Brookhaven National Laboratory (BNL). We describe how hardware and software configuration and the resource requirements determine the design of a batch system to be used for massive data reconstruction in High Energy Physics experiments.

INTRODUCTION

The Rhic Computing facility (RCF) is a large scale data processing facility at Brookhaven National Laboratory (BNL) serving the computing needs of the Relativistic Heavy Ion Collider (RHIC), a collider dedicated to high energy nuclear physics experiments.

The main source of CPU within RCF is the Linux farm, currently consisting of 1350 servers, divided among four RHIC experiments and RCF. Part of the nodes serve as data analysis farm operating LSF [1] batch system, while the rest is used for data reconstruction – namely reprocessing data collected by data acquisition systems of RHIC experiments and storing them on tapes, fur future analysis. The data reconstruction farm uses CRS, a home written batch system, written for the explicit purpose of managing RHIC data reconstruction jobs.

The mass storage for data files is provided by NFS servers and HPSS tape storage.

The CRS software for managing data reconstruction jobs does not scale with the size of the farm and as a result, as the farm grew in size the failure rate for CRS jobs became to high. As a result, the need for a new CRS software became evident.

REQUIREMENTS FOR THE NEW CRS

The new CRS system should satisfy several requirements:

- It should be highly reliable, well suited for mass data processing,
- It should be interfaced to HPSS tape servers and be able to stage files both to and from it. Staging data files should be done independently from further data processing, in order to optimize the operation to minimize the number of tape mounts.
- The system should allow operators to run data reconstruction jobs in a semi automatic way, the CRS system should be capable to identify failed jobs and, when possible, perform error diagnostics in order to establish cause of failure and advise the operator whether the error is recoverable or whether the job should be abandoned.
- The system should be able to keep track of the production by storing information about each job status and history in databases
- The system should provide users with possibility to check the progress of each individual job while it is in execution.

The requirements listed above suggested the Condor [2] batch system as the scheduler of the new CRS software. Condor provides users with possibility of defining jobs consisting of graphs of interdependent subjobs, which allow splitting the data staging operations from data processing, which in turns allows to make optimal use of the HPSS storage system.

The CRS system is interfaced to MySQL [3] databases which are used to record information about each job and subjob in the system, known input and output files – with their locations and statuses, as well as available network connections for moving data within the farm.

The user interface is provided by a GUI panel as well as by a set by line-mode commands.

CRS JOB

As first stem user describes the requested job by providing a file containing job specifications. These are written in a previously defined format and specify the job executables, input and output files as well as instructions for the system specifying definition of a "successful" job. (Namely: what should be the expected exit code, which of the expected output files must be present to consider the job to be successful, and which are optional).

Based on the user's requirements CRS writes scripts and job description files for the Condor batch system, registers the information about the job in databases and then submits it to the Condor Batch system.
Condor executes first subjobs which are responsible for staging the input data. Depending on the data type, the input files are located on NFS disks or in HPSS storage, and – in necessary – moved from tapes to HPSS cache disks.
Once all data staging subjobs are completed the main subjob starts data reconstruction. This step may take – depending on the experiment – from around half an hour to several hours.

After data reconstruction is over, the output files are checked for consistency and – if no errors are detected – they are exported to destinations requested by users - either NFS storage or HPSS tapes. In case of errors – which can be caused by factors ranging from bad data to mechanical breakdowns of the farm elements (disks, tape drives,..) the system tries to establish the cause of the problem, and take appropriate action. (Retry automatically to export data after faulty drive has been fixed, restart the job, or – in case of severe breakdown – halt the production and alert the operators).

SUMMARY

The RCF data reconstruction farm is currently replacing transition replacing its old, home written CRS batch system by a new, Condor based one. At the time of writing, one of the four RHIC experiments switched entirely to the new system, while the others are in process of changing to the new one. It is expected that replacing the old CRS by a new. Condor based one, will lead to more efficient usage of the RCF computing resources.

ACKNOWLEDGEMENTS

The RCF would like to thank BNL's Physics Department, Information Technology Division (ITD) and the funding agencies (DOE,NSF, etc) for their support.

REFERENCES

[1] http://platform.com

[2] http://www.cs.wisc.edu/condor

[3] http://www.mysql.com/

---

* tomw@bnl.gov

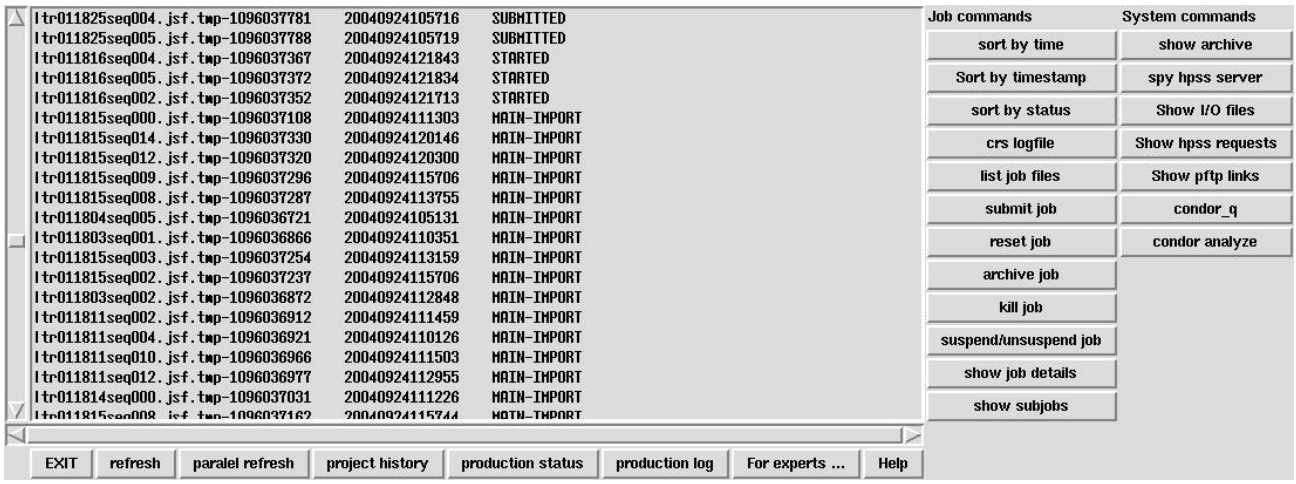

Figure 1: Main control panel of the CRS batch system.