

Deployment of JIM for the CDF Grid

Mòrag Burgon-Lyon^a, Gabriele Garzoglio^b, Valeria Bartsch^c, Richard St.Denis^a, Stefan Stonjek^c

^aUniversity of Glasgow, Scotland, United Kingdom

^bComputing Division, Fermi National Accelerator Laboratory, Batavia, IL, USA

^cUniversity of Oxford, England, United Kingdom, based at Fermi National Accelerator Laboratory

The CDF (Collider Detector at Fermilab) Experiment[1] is in the process of distributing computing infrastructure to numerous sites world-wide. The software used for this task is comprised of: a mature data handling system called SAM (Sequential Access to data via Metadata)[2]; DCAF (Decentralised CDF Analysis Farm)[3] for local job queuing and execution; and JIM (Job and Information Management)[4] is used to collect and distribute jobs to SAM stations and DCAF farms. These three core elements of CDF Grid are introduced and the plans, problems and solutions in deploying the JIM component are discussed.

1. Introduction

The CDF High Energy Physics Experiment has declared the intention to have 25% of computing CPU from external resources by June 2004 and 50% by June 2005. To achieve this target SAM and DCAF have been installed at a number of sites. At the time of writing the June 2004 target has been met, as there are currently five production sites totaling 850GHz of available CPU in addition to two sites for testing and several sites under development.

This combination of software allows users to submit jobs choosing the site where they will be executed. SAM handles the delivery of files and DCAF executes the job returning the output sandbox to the user in the form of a tar file. This existing pool of resources is constantly being expanded.

The next step is a grid solution whereby a user submits a job and it is automatically sent to the most appropriate site, usually where most or all of the required data is present. Depending on the location of the required data, and resource availability JIM will send the job to the best suited site rather than strictly taking the job to the data or the data to the job. To achieve this, an appropriate remote JIM installation will be added to each site providing the grid functionality.

2. Job and Information Management

JIM has been written as a grid extension to SAM. JIM and SAM together are known as SAM-Grid. JIM can be thought of as having two parts, Job Handling and Job Monitoring[5]. These can be further broken down into four installation components: a client site for job submission; submission site for job queuing; monitoring site keeps track of resources, services and jobs; an execution site advertises resources and services, and runs the job.

The thin client software can be installed on any PC and is configured to send jobs to a single submission site, with one submission site providing services to a large physical area, eg. Scotland. The submission site currently uses the Condor match making service to automatically select the job execution location. A site's resources are described using XML with subsequent projections onto the Condor Classified Advertisements (ClassAd) framework. Condor has been extended such that it communicates dataset distribution information with SAM to influence the choice of execution site[6]. SAM can be used to transfer large volumes of data from any SAM cache to the execution site, and also to store files back into the main SAM database.

Through the client site software, JIM presents

the user with uniformity of submission, regardless of where the job is submitted or actually runs. From a site management perspective there is a generic infrastructure with plug-ins to separate services.

When development first began on the JIM extension to SAM, the grid tools available were too limited for the task. Initially the JIM developers extended the Condor match making service to provide sufficient functionality to allow JIM to operate with SAM as a full grid service. As grid technologies have advanced, the JIM development team have made changes to use these new and improved technologies. An example of this has been the updating of SAM-Grid software for compatibility with Globus 2.4.3.

The Virtual Data Toolkit (VDT)[7] is an ensemble of grid middleware that can be easily installed and configured. VDT supports the LHC Computing Grid (LCG) and the Particle Physics Data Grid (PPDG). Distribution of SAM-Grid is now available using VDT.

3. Implementation Issues

In recent months the JIM development team has focused on Monte Carlo (MC) production for the D0 experiment. The D0 success rate for MC is now over 99%. Efforts are underway to emulate these results for CDF. A script that makes a tarball from the CDF software environment has been used to manually run CDF MC on D0 computing facilities at Wisconsin. This first step means that the CDF software environment can be transferred around the grid, preventing problems with differing code versions on execution sites. This ensures that shared resources can be used fully for CDF jobs without application version issues.

MC can be split into three separate parts, each taking the output from the previous step as input. When a job is first submitted, a list of run numbers is given. The run number list is filtered using a database query to ensure only runs for which the luminosity makes sense will be used. The output is passed through simulation, reconstruction and production.

The MC methods differ between the two experiments, thus extra development effort is in place

to change both JIM and CDF MC code. This will allow the collation of information from earlier parts of the MC generation to ensure that any failed jobs will not adversely effect the end results. Resubmissions of failed jobs can be generated using an appropriate random number seed.

SAM station installations have been vastly simplified by the creation of a script. The once timely and difficult process can now be completed within a couple of hours, largely unattended. Simplifying the installation procedure was a critical step to allow the quick roll-out of SAM, a crucial element of the CDF Grid software setup. A similar script is currently under development to provide the same ease of installation for JIM.

The JIM installation script is slightly more complex than the SAM script due to the fact there are four separate types of installation, giving several possible combinations of site. In practise however, only three types of installation are common (client only; execution and monitoring; and all four at submission sites). As JIM can interface to many different batch systems this also presents problems as the this information requires knowledge of the local system, and cannot be readily scripted[9].

CDF uses local databases to handle product installations, avoiding the need for root access. The installations of SAM and JIM software are segregated into two separate databases. The SAM Station script is presently used regularly by participating sites to update SAM installations, keeping them current. Unfortunately some of the JIM configuration is required in the SAM products database. With the inclusion of JIM on these sites, the SAM Station script requires a rewrite to backup and then restore this configuration on upgrading SAM. A manual procedure for updating will be used until the SAM installation script can be amended.

4. Deployment Plans

The following diagram, figure 1 shows how elements of the CDF Grid fit together. Users currently submit jobs from their terminal to DCAF. Although this will still be an option, once JIM is deployed, users will be encouraged to submit

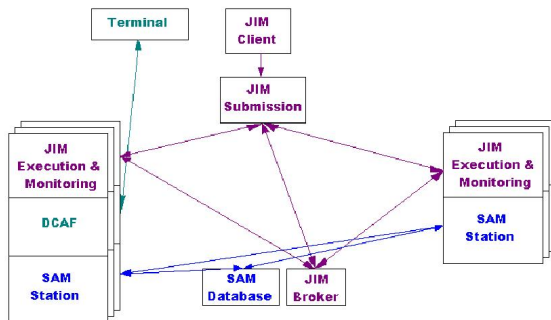


Figure 1. How CDF Grid fits together

their jobs through the JIM client software. JIM client passes the job to the area submission site for queueing. After communicating with the broker, the job will be sent to an execution site, which may have a DCAF. The job will be executed, using SAM to transfer files, and DCAF or the local batch system (e.g. PBS) to execute the job.

JIM deployment is planned for existing SAM stations using the JIM installation script. The script is currently being tested and refined using a small test cluster based at Fermilab. Installations of the latest version of SAM-Grid are planned for sites at the University of Glasgow, using a PC frontend to ScotGrid[8], and an eleven node Dell cluster at the University of Oxford.

The steps to full CDF grid deployment entail testing a complete JIM and SAM installation at the Universities of Glasgow and Oxford; addition of JIM to a site running DCAF, and finally roll-out of JIM software using the JIM installation script. Testing will focus on Monte Carlo and Analysis jobs and will also include Reconstruction as these are the main job types used for the experiment.

The deployment will be separated into physical regions (e.g. Scotland, Southern England) where each physical region will have one submission site installation. Locations wishing to use the JIM interface to CDF Grid will require a JIM client installation. Local installations of the small JIM

client software will be configured to pass submitted jobs to the regional submission site. Monitoring site installations will be required at execution sites. Each site that runs DCAF and SAM will be expanded with at least the JIM execution site installation. As some sites will require all four installation types, the JIM installation script will allow the local system administrator to pass over installations the site does not require.

Due to the issues presented in section 3, the planned deployment of JIM software has slipped. As these issues are being addressed, the deployment is now planned for the Autumn.

5. Summary

CDF have sufficient DCAF installations to allow 25% of computing to be done offsite. To achieve the goal of 50% by June 2005 SAM-Grid is required, as resources where DCAF installations are not possible may be harnessed.

JIM deployment has been delayed while MC issues are being resolved, and an installation script is developed and tested. Progress has been made, and it is envisioned that deployment will be underway by Autumn.

REFERENCES

1. CDF home page
<http://www-cdf.fnal.gov/>
2. SAM home page
<http://fcdfsun1-prd.fnal.gov/sam/>
3. DCAF home page
<http://cdfcaf.fnal.gov/>
4. JIM home page
<http://www-d0.fnal.gov/computing/grid/>
5. M. Burgon-Lyon, *et al.*, Experience using grid tools for CDF Physics, Advanced Computing and Analysis Techniques in Physics Research (ACAT) 2003.
6. A. Baranovski, *et al.*, Management of Grid Jobs and Information within SAMGrid, Computing for High Energy Physics (CHEP) 2003, La Jolla, CA.
7. Virtual Data Toolkit
<http://www.cs.wisc.edu/vdt/>

8. ScotGrid home page
<http://www.scotgrid.ac.uk/>
9. G. Garzoglio, I. Terekhov, A. Baranovski, S. Veseli, L. Lueking, P. Mhashikar, V. Murthi, The SAM-Grid Fabric Services, Advanced Computing and Analysis Techniques in Physics Research (ACAT) 2003.