

grid-control - The Swiss Army knife of job submission tools

Tuesday, 11 October 2016 16:30 (15 minutes)

grid-control is an open source job submission tool that supports common HEP workflows. Since 2007 it has been used by a number of HEP analyses to process tasks which routinely reach the order of tens of thousands of jobs.

The tool is very easy to deploy, either from its repository or the python package index (pypi). The project aims at being lightweight and portable. It can run in virtually any environment with access to some submission infrastructure. To achieve this, it avoids any external dependencies. Only a python interpreter (CPython or pypy) supporting any language version from 2.3 upwards, including 3.x, is required.

The program supports job submission to a wide range of local batch systems and grid middleware. For small tasks and tests, it is also possible to run jobs on machines without any batch system.

grid-control is built around a powerful plugin and configuration system, that allows to easily define the workflow that is to be processed.

A particularly useful feature for HEP applications is the job parameter system built into grid-control. It provides a convenient way to define the parameter space that a task is based on. This parameter space can be build up from any number of variables and data sources, which can have complex dependencies on each other. The parameter system is able to handle changes to the data source as well as to other parameters. It can transparently adapt the job submission to the new parameter space at runtime.

grid-control provides several plugins to retrieve datasets from different sources. This ranges from simple file listings, to databases, directory contents, other grid-control tasks and more. These datasets may contain a list of URLs, an optional number of work units (eg. events or file size), arbitrary metadata (eg. lumi section information or generator parameters) and locality information.

All datasets are processed through a configurable pipeline of dataset filters, partition plugins and partition filters. Several methods to split datasets into partitions are supported. These partition plugins can take the number of files, size of the work units, metadata or combinations thereof into account.

Dataset changes on the file level (additions or removals) as well as on the work unit level (expanding or shrinking files), are propagated through the processing pipeline and transparently trigger adjustments to the processed parameter space.

For HEP workflows this allows to run jobs on an ever expanding dataset with a single processing task that regularly queries some datasource and spawns new jobs.

While the core functionality is completely experiment independent, a dedicated plugin is available to simplify running tasks using the CMS experiment software (version 1.x-8.x) and to use the CMS Dataset Bookkeeping Service, CMS Data Aggregation System or PhEDEx as data source.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Distributed workload management

Primary Keyword (Mandatory)

Data processing workflows and frameworks/pipelines

Primary author: Dr STOBER, Fred (Hamburg University (DE))

Co-author: FISCHER, Max (KIT - Karlsruhe Institute of Technology (DE))

Presenter: Dr STOBER, Fred (Hamburg University (DE))

Session Classification: Posters A / Break

Track Classification: Track 7: Middleware, Monitoring and Accounting