Abstract

Current SLURM versions are featured with a variety of job completion logging plugins such as *mysql*, *filetxt* or *script*. The purpose of this presentation is to raise awareness about a new *jobcomp-family* related plugin that manages the job completion information so that it is stored into a user configured *elasticsearch*[1] cluster. An elasticsearch cluster structure is composed by nodes, indexes, types and documents. A document is a basic unit of information that can be indexed, and it is expressed in JSON[2]. Elasticsearch operations are performed through HTTP requests against a configured URL in the *JobCompLoc* parameter inside the SLURM configuration file. The main functionalities of this new plugin include checking whether the configured elasticsearch server is accessible and operational, preprocess the job completion data so that it is encoded in a proper JSON format and saving state of the data that could not be indexed (for instance because the elasticsearch index is temporarily in *read-only* mode) in the *StateSaveLocation* so that it can be restored for further indexing retries. Plugin implementation has been held in a context that is part of a bigger project for a Final Master Thesis in the Barcelona School of Informatics[3] and in collaboration with the Barcelona Supercomputing Center[4].

[1] Elasticsearch is a highly scalable open-source full-text search and analytics engine. It allows you to store, search, and analyze big volumes of data quickly and in near real time: `http://www.elasticsearch.org/`.

[2] JavaScript Object Notation, which is an ubiquitous internet data interchange format.

[3] Facultat d'Informàtica de Barcelona: `http://www.fib.upc.edu/en.html`

[4] Barcelona Supercomputing Center: `http://www.bsc.es/`