

Slurm workbench: a cluster visualized research system based on Slurm REST API

Tuesday, 22 March 2022 13:30 (30 minutes)

As many MPI and GPU computing requirements are raised from experiments, the computing center of IHEP founded the Slurm cluster in 2017 and started the parallel computing services afterwards. Since then, users and applications of the Slurm cluster are increased from time to time. By the end of 2021, there are 16 applications and 160 users served by more than 6200 CPU cores and 200 GPU cards from the Slurm cluster.

Slurm provides command lines to submit jobs, query and control cluster status. Those commands are powerful and comprehensive. However, from the view of administrator, a well-functional cluster not only asks for command lines, but also visualized support systems. Visualized systems can help administrators to monitor real-time cluster status, generate statistics from historic job records, and submit specific pattern of jobs out of research purpose. Such visualized systems are formed to be the Slurm ecosystem on top of the cluster itself.

Slurm started to release REST APIs since version 20.02. Slurm REST APIs can be used to interact with slurmctld and slurmd daemons, so that job submission and cluster management can be achieved on a web interface directly. In addition, responses from slurmctld and slurmd in JSON could be organized properly in a favored way by cluster administrators.

This paper presents the slurm workbench system, which is developed with Python Flask based on Slurm REST APIs. Slurm workbench is consisted with three subsystems, which are dashboard, jasmine and cosmos. Dashboard is developed to display current cluster status including jobs and nodes. Jasmine is used to generate and submit specific pattern of jobs according to job parameters, which is convenient to study resource allocation and job scheduling. Cosmos is a job accounting and analysis system, with which job statistical charts are generated based on history job records. With jasmine, cosmos and dashboard working together, Slurm workbench provides a visualized way to study application and Slurm cluster.

Primary author: DU, Ran (Institute of High Energy Physics, Chinese Academy of Sciences)

Co-authors: SHI, Jingyan (IHEP); JIANG, Xiaowei (Institute of High Energy Physics, Chinese Academy of Sciences)

Presenter: DU, Ran (Institute of High Energy Physics, Chinese Academy of Sciences)

Session Classification: Network, Security, Infrastructure & Operations

Track Classification: Track 7: Network, Security, Infrastructure & Operations