

Frank Winkler

Slurm PrEp Plugin

NHR Monitoring Workshop 2022
07/18/2022 to 07/19/2022, Erlangen

Slurm PrEp Plugin

- Provides a native C interface to the same hooks used by the **Prolog**, **Epilog**, **PrologSlurmctld**, and **EpilogSlurmctld** scripts
- Template for a PrEp plugin

```
const char plugin_name[] = "PrEp plugin template";
const char plugin_type[] = "prep/template";
const uint32_t plugin_version = SLURM_VERSION_NUMBER;

int init(void)
void fini(void)

extern int prep_p_prolog_slurmctld(job_record_t *job_ptr, bool *async) {
    ...
}
extern int prep_p_epilog_slurmctld(job_record_t *job_ptr, bool *async) {
    ...
}
```

- See: https://slurm.schedmd.com/prep_plugins.html

Slurm PrEp Plugin

- The following metadata can be acquired at prolog and epilog:
 - Unique job identifier, ArrayID
 - Project, user, job name
 - Start and end time, walltime
 - Status (running, completed, timeout, failed, OOM, cancelled)
 - Partition
 - Allocated compute nodes
 - Allocated CPUs/GPUs on each node
 - Exclusive nodes
 - Main memory
 - Working directory
 - Job script

Example: https://gitlab.hrz.tu-chemnitz.de/pika/pika-packages/-/blob/master/pika-prep-plugin/slurm-prep-pika_v4.c

Slurm PrEp Plugin

Build and Install:

- Build plugin as shared library
- Update slurm.conf → PrEpPlugins=[plugin name]
- Restart slurm controller

Pro:

- Centralized instance on the Slurm controller that can write all required data into the database
- Message broker enables multiple clients to capture specific job data

Contra:

- Not necessarily compatible with newer Slurm versions
- For newer Slurm versions small changes have to be made and tested