

HPC working seminar for physicists

Scientific Computing Department at HIM

Dr. Dalibor Djukanovic

Dr. Peter-Bernd Otte

bi-weekly meeting – 29.3.2022



Today's Topics

1. News
 2. SLURM from inside Container
 3. your questions / discussion / requests to the maintainers
- compact in time (15mins + user questions/discussion).
 - bring people together tackling the same problems
 - minutes: <https://www.hi-mainz.de/research/computing/hpc-working-seminar/>

Hot Topics we are working on

- Theory partition:
 - + 900TB, quota updates this afternoon (15.3.)
- Lustre stripe settings:
 - Stripe count = 4, threshold > ~1000 MB
 - For miifs01, not for fs04 and fs05

```
[[pbotte@login23 scratch]$ lfs getstripe -d /lustre/miifs01
```

```
lcm_layout_gen: 0
lcm_mirror_count: 1
lcm_entry_count: 2
lcme_id: N/A
lcme_mirror_id: N/A
lcme_flags: 0
lcme_extent.e_start: 0
lcme_extent.e_end: 1073741824
  stripe_count: 1      stripe_size: 1048576      pattern:      raid0      stripe_offset: -1
lcme_id: N/A
lcme_mirror_id: N/A
lcme_flags: 0
lcme_extent.e_start: 1073741824
lcme_extent.e_end: EOF
  stripe_count: 4      stripe_size: 1048576      pattern:      raid0      stripe_offset: -1
```

- Remember:
 - 1k files / dir
 - prefer Large files!

Baby sitting of jobs

- Jobs sometimes fail, due to different reasons
 - baby sitting necessary
- problematic, if $N \gg 1$

- Solution: a script takes over the task

Baby sitting of jobs

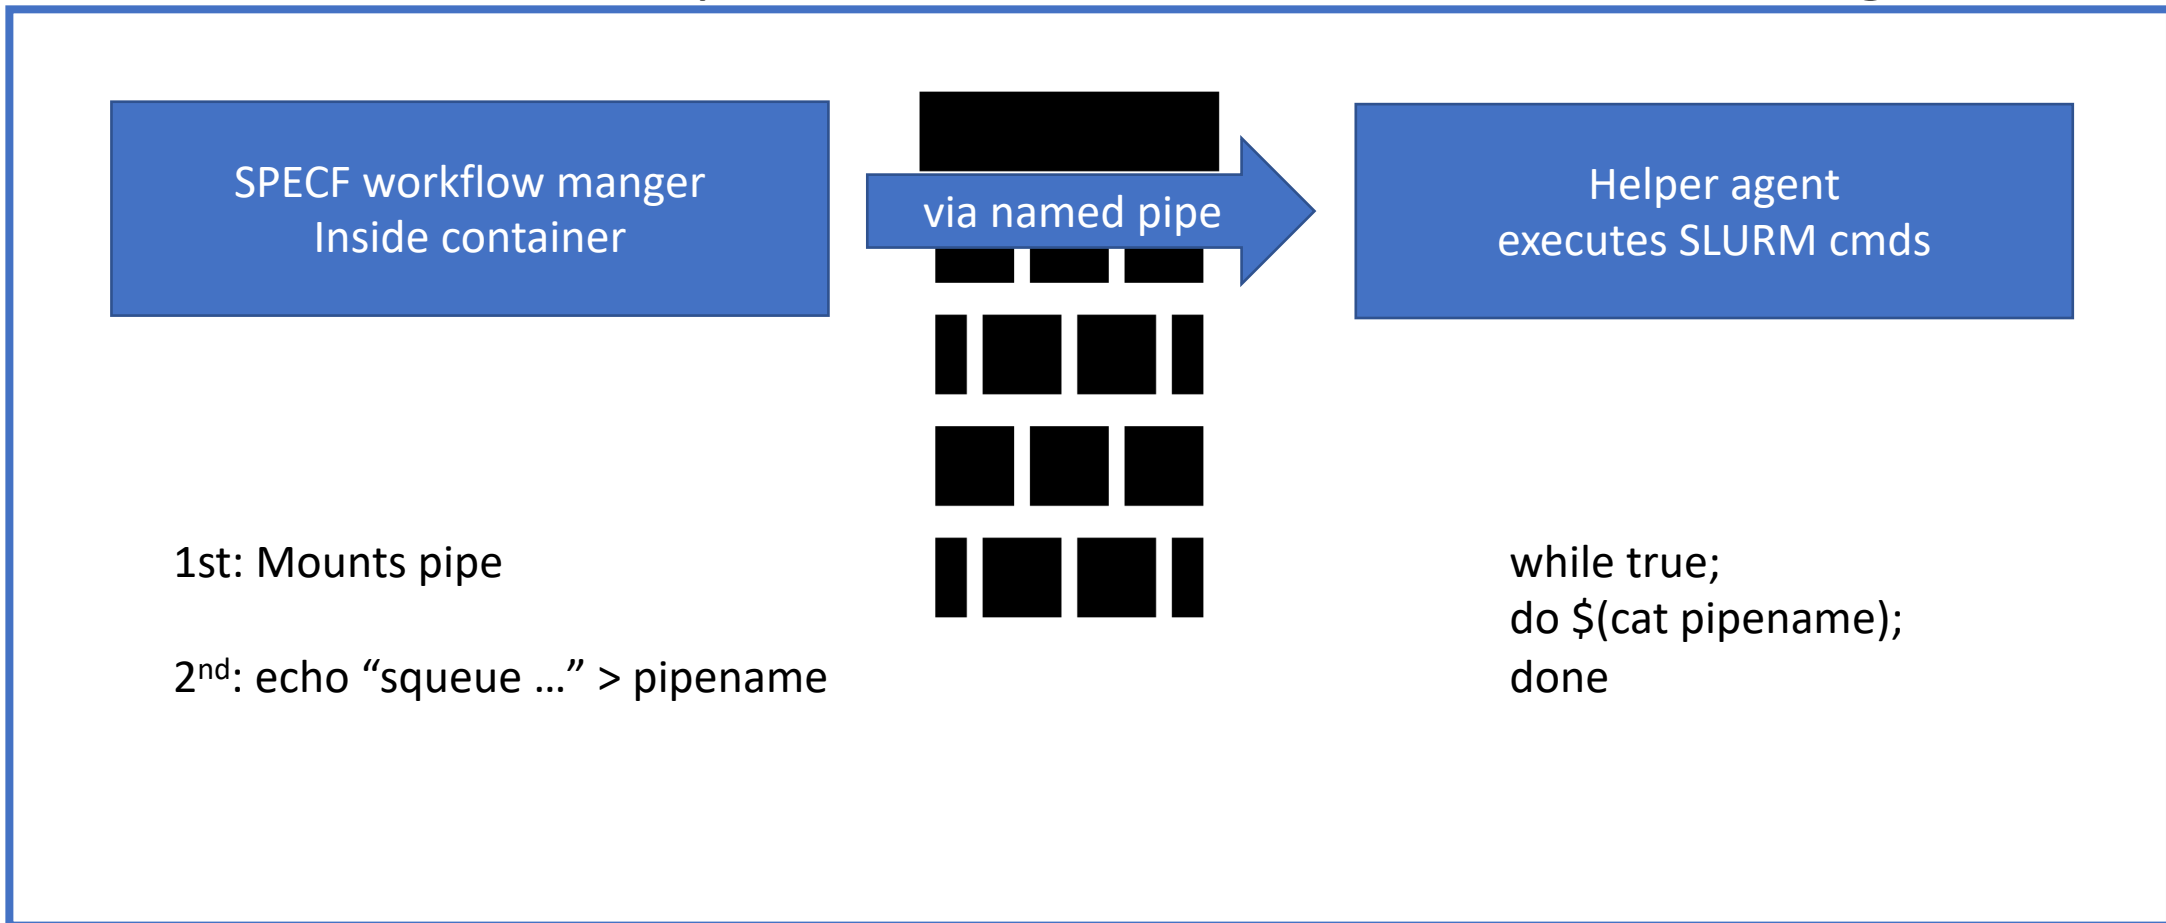
- What if $N \gg 1$ and baby sitting is complex?
 - eg 20k jobs for energy or angle scan
 - different calibrations
 - multiple analysis steps
- Solution: a decent baby sitting app
 - “workflow manager”
- SPECIF now came up with a containerised workflow manager

Workflow manager for jobs

- SPECIFIC now came up with a containerised workflow manager
 - <https://github.com/panda-luminosity-detector-group/LuminosityFit/>
- Problem on Himster 2: SLURM commands not accessible from inside container
 - SLURM setup changes often and high complexity
 - SLURM sits in standard directory, so no mount possible
- Solution: a helper agent outside the container

Workflow manager for jobs

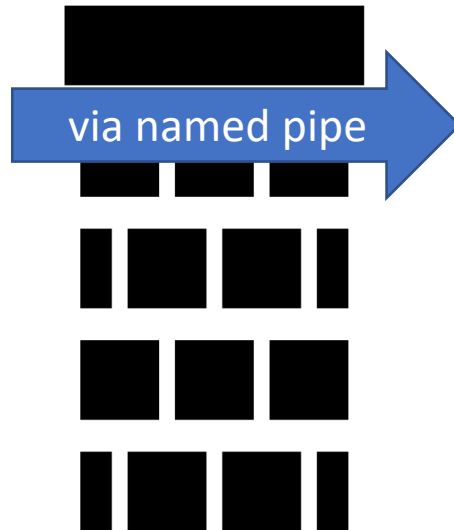
- SPECf now came up with a containerised workflow manager



Workflow manager for jobs

- SPECf now came up with a containerised workflow manager

SPECf workflow manager
Inside container



Helper agent
executes SLURM cmds

1st: Mounts pipe

2nd: echo "squeue ..." > pipename

Python script

```
Executable File: 227 lines (188 sloc) | 7.23 KB
1 #!/usr/bin/env python3
2 """
3 Simple program that runs commands given as a json-formatted object
4 in a named pipe and returns their outputs and return code in the same way.
5
6 the structure of the order and return object is governed by the SlurmOrder Class.
7 """
8
9 import attr
10 import datetime
11 import json
12 import logging
13 import os
14 import shlex
15 import subprocess
16 import stat
17 import sys
18 from typing import Dict
19
20
21 @attr.s(hash=True)
22 class SlurmOrder:
23     orderType: str = attr.ib(default="regular")
24     cmd: str = attr.ib(default="id")
25     returnCode: int = attr.ib(default=-1)
```

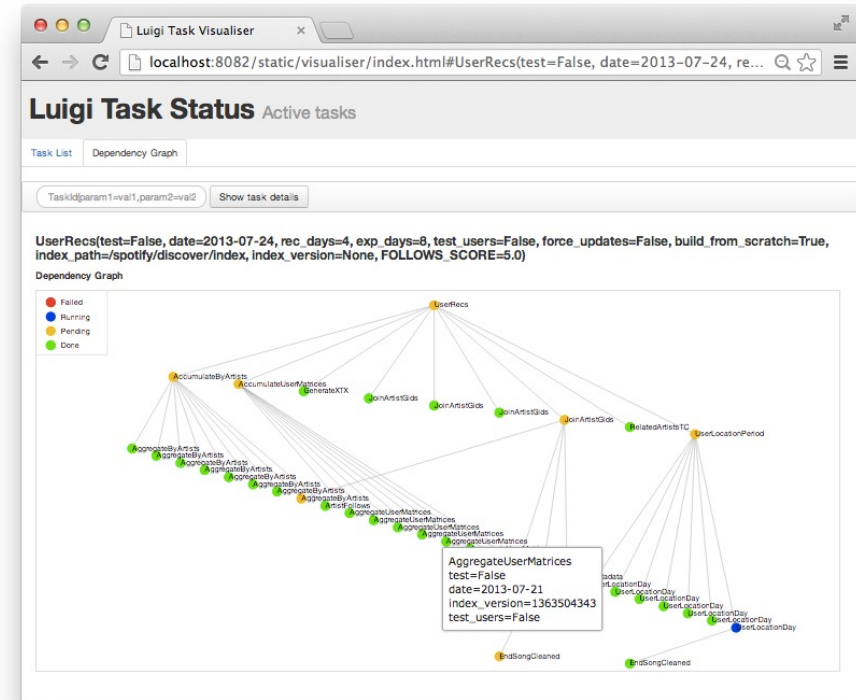

Alternative Workflow manager



<https://github.com/spotify/luigi>

pip install luigi

A screenshot of the Luigi Task Status web interface. The top navigation bar includes "Task List", "Dependency Graph", "Workers", and "Resources". A summary section shows: Pending Tasks: 3, Running Tasks: 1, Batch Running Tasks: 0, Done Tasks: 134, Failed Tasks: 3, Upstream Failure: 1, Disabled Tasks: 0, and Upstream Disabled: 0. A table below lists tasks, with one task "LoadTable" highlighted as running. The table has columns for Name, Details, Priority, Time, and Actions.

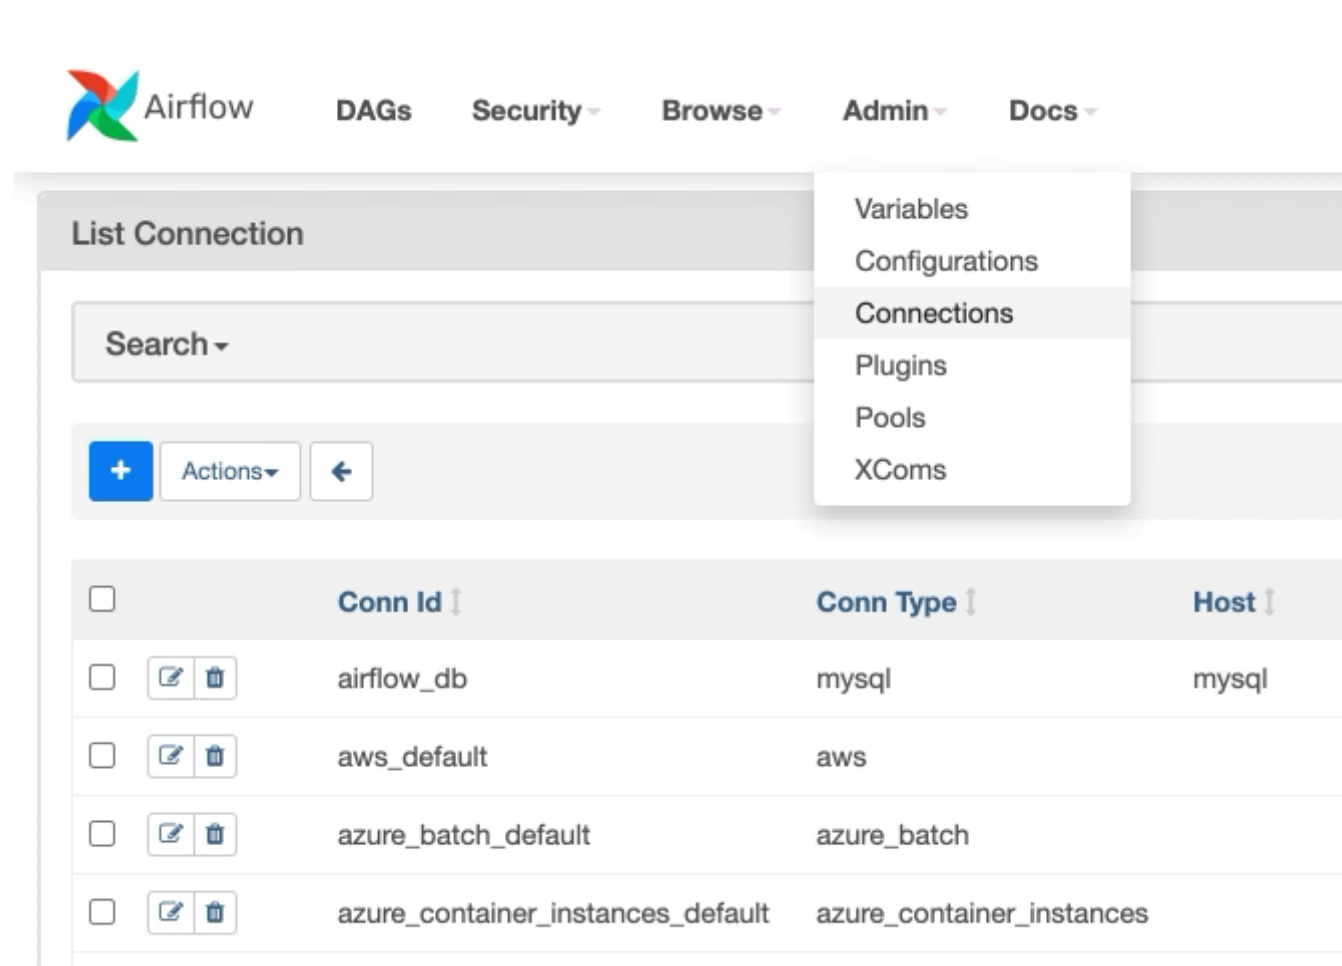


Dependency graph manger

Alternative Workflow manager

- Apache Airflow

<https://airflow.apache.org>



The screenshot displays the Apache Airflow web interface. At the top, the Airflow logo is followed by navigation links: DAGs, Security, Browse, Admin, and Docs. The 'Admin' dropdown menu is open, showing options: Variables, Configurations, Connections, Plugins, Pools, and XComs. The main content area is titled 'List Connection' and features a search bar, a '+ Actions' button, and a table of connections.

<input type="checkbox"/>	Conn Id ↓	Conn Type ↓	Host ↓
<input type="checkbox"/>	airflow_db	mysql	mysql
<input type="checkbox"/>	aws_default	aws	
<input type="checkbox"/>	azure_batch_default	azure_batch	
<input type="checkbox"/>	azure_container_instances_default	azure_container_instances	

your questions / discussion /
requests to the maintainers?

Next meeting, 12th April 2022

- Only technical / short meetings during semester break
- hand in your topics!