

HEPIX Autumn 2024 (Oklahoma)

Credits: J. Flix / K. Fàbrega

Date: 7/11/2024

```
In [1]: import numpy as np
import pandas as pd
import os
import matplotlib as mpl
import matplotlib.pyplot as plt
from time import strftime, localtime
import sys
import glob
import re
```

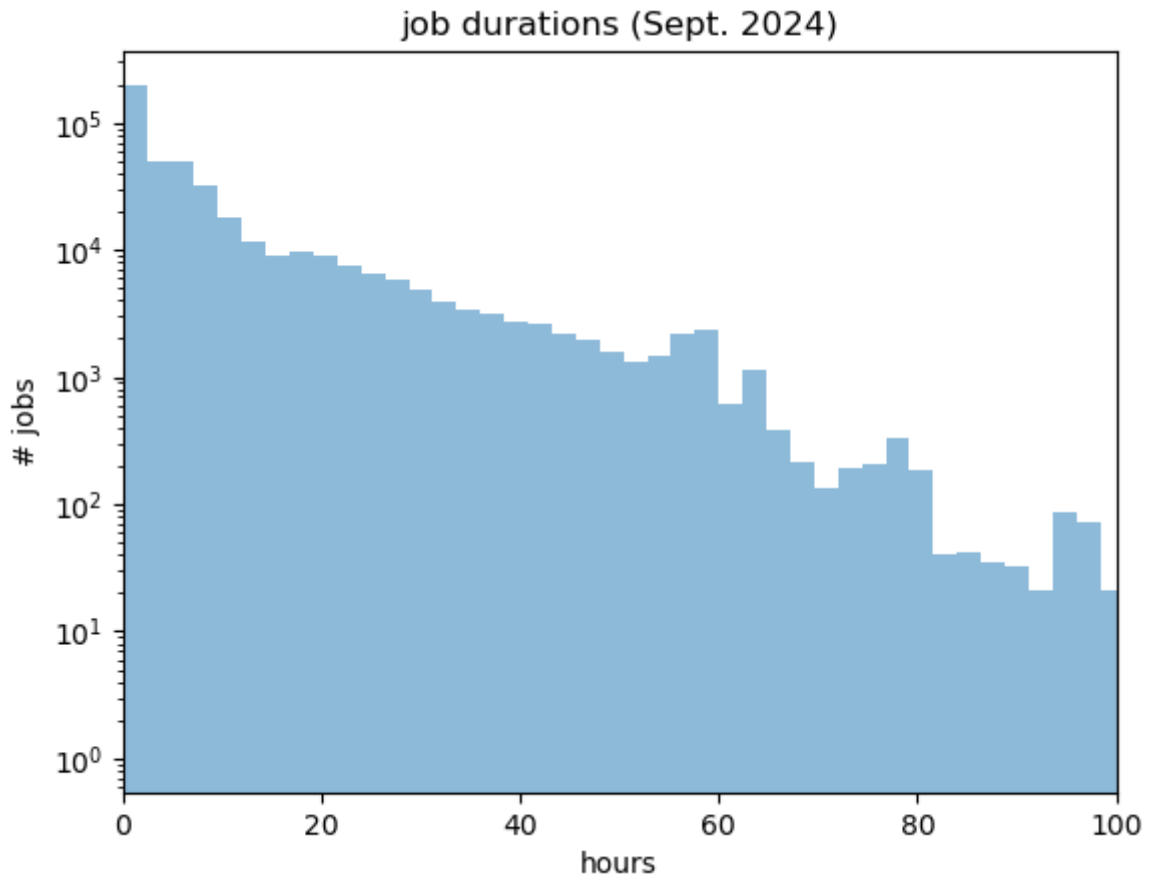
Read September 2024 data from 5 HTCondor-CE's + local submits at PIC

```
In [2]: df = pd.read_pickle("./df_sept2024.pkl")
df.sort_values('timestamp', ascending=True, inplace=True)
```

Plot job durations

The PIC local HTCondor config allows for jobs to extend up to 100h

```
In [3]: plt.hist(df['jobduration']/3600, bins=100, log=True, alpha=0.5)
plt.xlabel("hours")
plt.ylabel("# jobs")
plt.title("job durations (Sept. 2024)")
plt.xlim(0,100)
plt.show()
```



Resemble PIC farm utilization

From the job timestamps (start_time and end_time), and the number of cores used per job, we can resemble the PIC farm utilization.

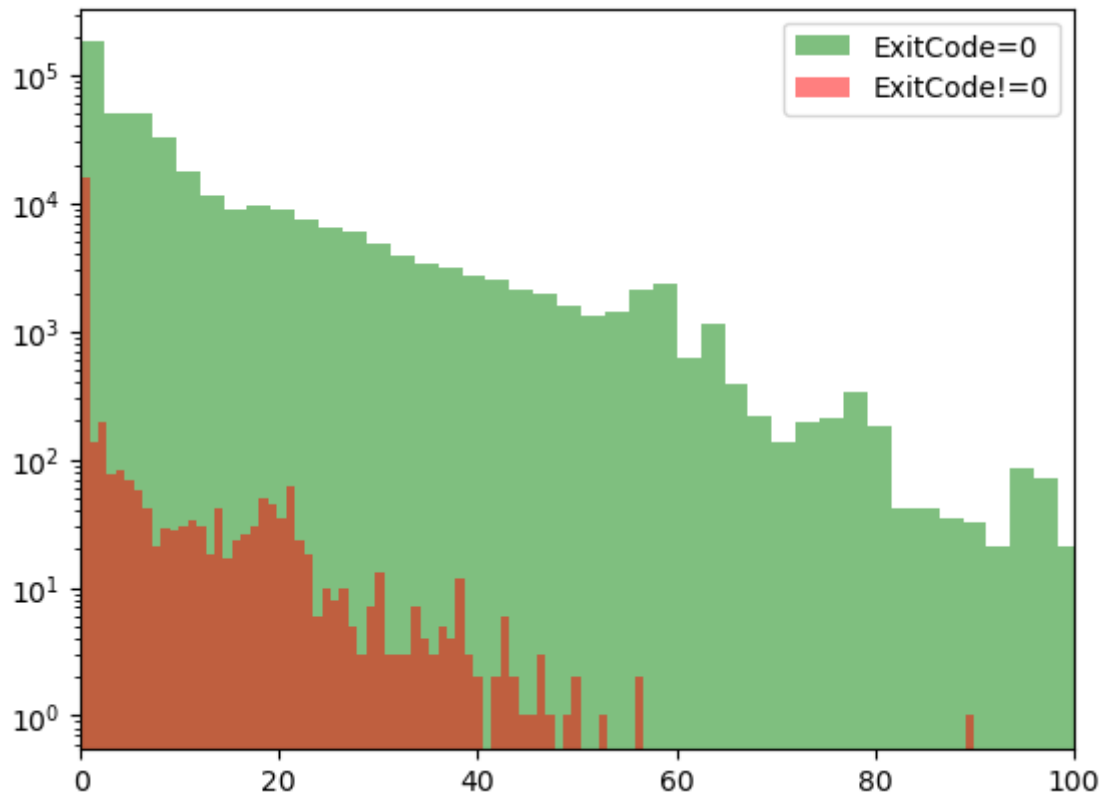
Querying in intervals of 60 seconds for the whole period, we can compute how many jobs were running at a particular time (i.e. 'startdate' <= t & 'enddate' > t). Of course, we miss those jobs that have a duration of < 60 seconds. There are ~6% of jobs failing there. Of course, the 60 seconds default can be lowered, but the computational time increase. We use 60 seconds for this live demo.

```
In [4]: print("# jobs:", df.shape[0])
print("# jobs (<60 s):", df[df['jobduration']<60].shape[0])
print("Fraction (%)ate", 100.*df[df['jobduration']<60].shape[0]/df.shape[0])

# jobs: 445455
# jobs (<60 s): 26094
Fraction (%): 5.857830757315554
```

Indeed the jobs we are losing are those that mostly fail...


```
In [5]: plt.hist(df['jobduration'][df['ExitCode'] == "0"]/3600,bins=100, color="g")
plt.hist(df['jobduration'][df['ExitCode'] != "0"]/3600,bins=100, color="r")
plt.xlim(0,100)
plt.legend()
plt.show()
```



Computation of running cores for September 2024 at PIC

```
In [6]: t_fi = []
v_fi = []
for ts in np.arange(int(df['enddate'].min()),int(df['enddate'].max()),60):
    t_fi.append(ts)
    v_fi.append(df['request_cpus'][(df['startdate'] <= ts) & (df['enddate']
    <= ts)])

plt.plot(t_fi, v_fi)
plt.ylabel("Running cores")
plt.show()
```

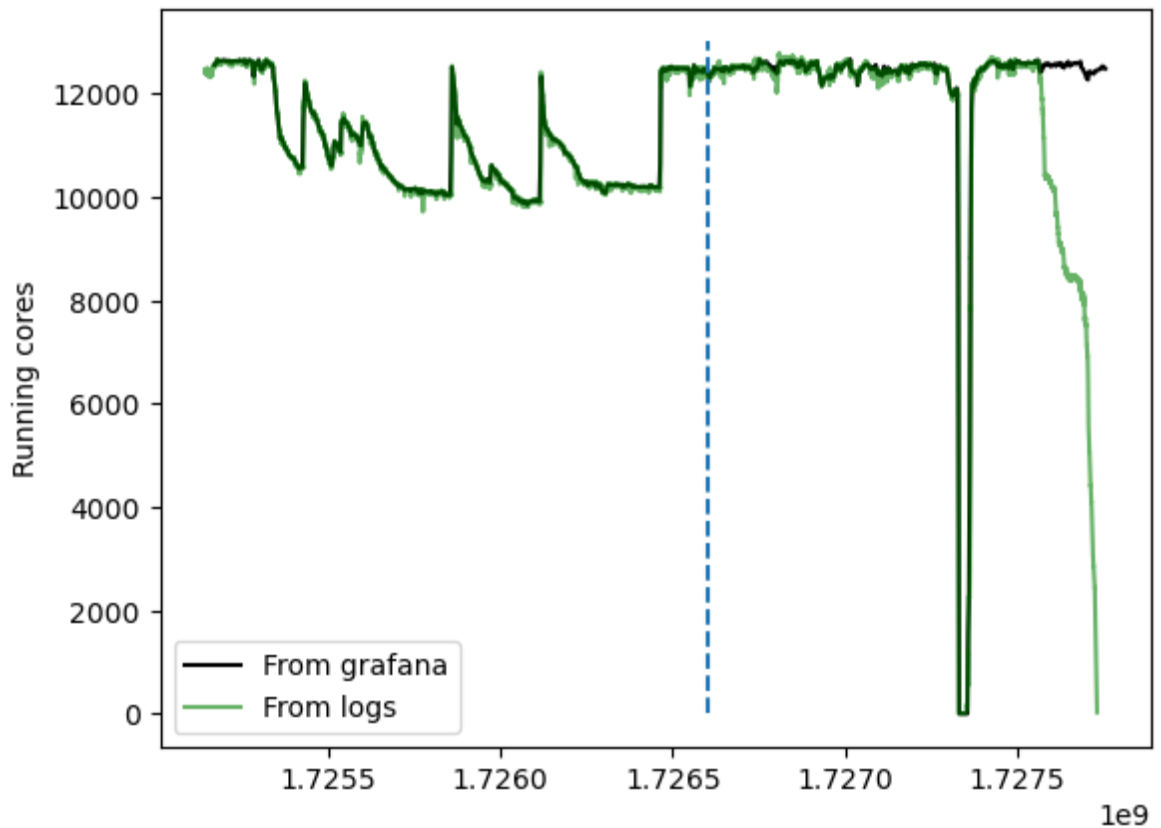


Compare the estimated running cores with the running cores we monitor in realtime in PIC Grafana for Sept. 2024. At the end of the period, of course, we cannot resemble the farm utilization because there are missing jobs (last 100h).

```
In [7]: time_drop=1726600000

df_graf = pd.read_csv("./HTCondor_Grafana_all.csv", sep=",", header=0)

plt.plot(df_graf['Time']/1000, df_graf['All'], label="From grafana", color="red", alpha=0.6)
plt.plot(t-fi, v-fi, label="From logs", color="green", alpha=0.6)
plt.vlines(time_drop, 0, 13000, linestyle="--")
plt.ylabel("Running cores")
plt.legend(loc='best')
plt.show()
```



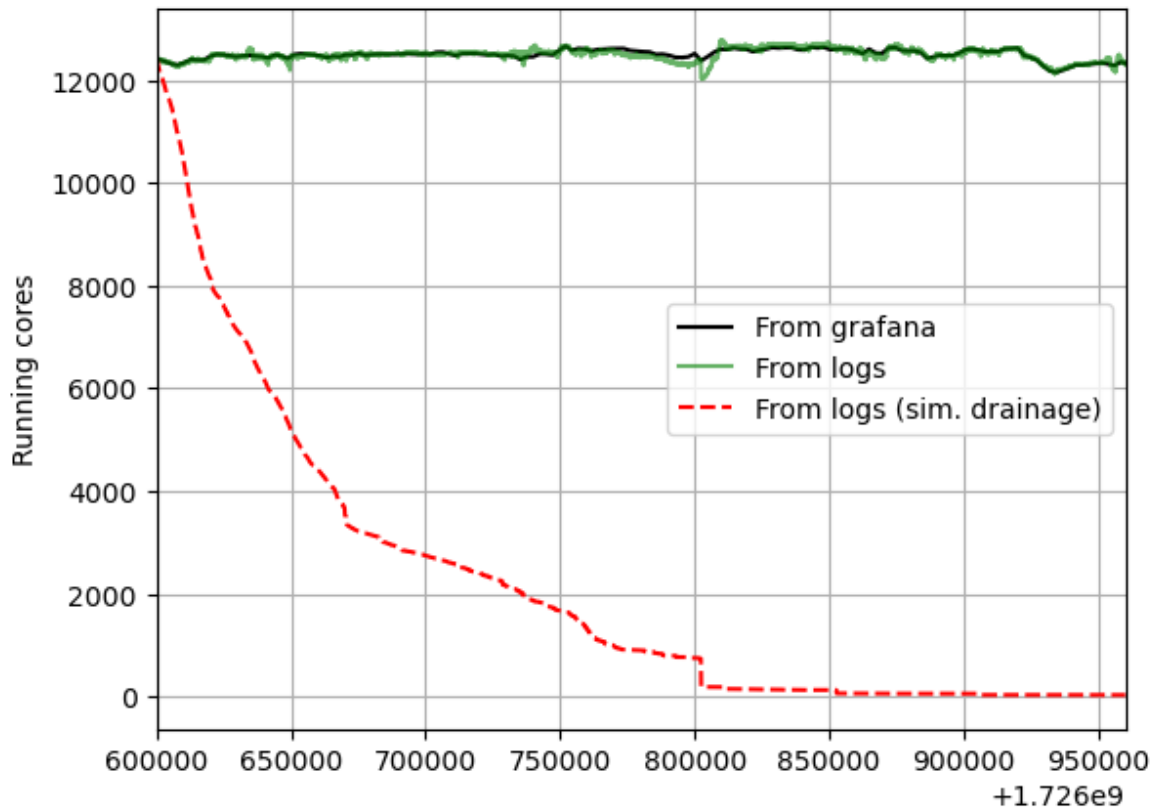
Simulating a natural job drainage

Defining a time in which we stop sending jobs to HTCondor (time_drop), it is easy to simulate a natural job drainage at PIC farm. Simply we need to drop all of the jobs in the dataframe that started after this time_drop (blue dashed line in previous plot).

```
In [8]: df_drop = df.drop(df[ df['startdate'] > time_drop].index)
```

```
In [9]: t_fi2 = []
v_fi2 = []
for ts in np.arange(int(df_drop['enddate'].min()),int(df_drop['enddate'].max()),int(df_drop['enddate'].max()-int(df_drop['enddate'].min()))):
    t_fi2.append(ts)
    v_fi2.append(df_drop['request_cpus'][(df_drop['startdate'] <= ts) & (df_drop['enddate'] > ts)])

plt.plot(df_graf['Time']/1000, df_graf['All'], label="From grafana", color="black", alpha=0.6)
plt.plot(t_fi, v_fi, label="From logs", color="green", alpha=0.6)
plt.plot(t_fi2, v_fi2, label="From logs (sim. drainage)", color="red", alpha=0.6)
plt.xlim(time_drop,time_drop+100*3600)
plt.ylabel("Running cores")
plt.legend(loc='best')
plt.grid()
plt.show()
```



One can see that the farm empties after 100h, which is in agreement to the job duration distribution we made earlier.

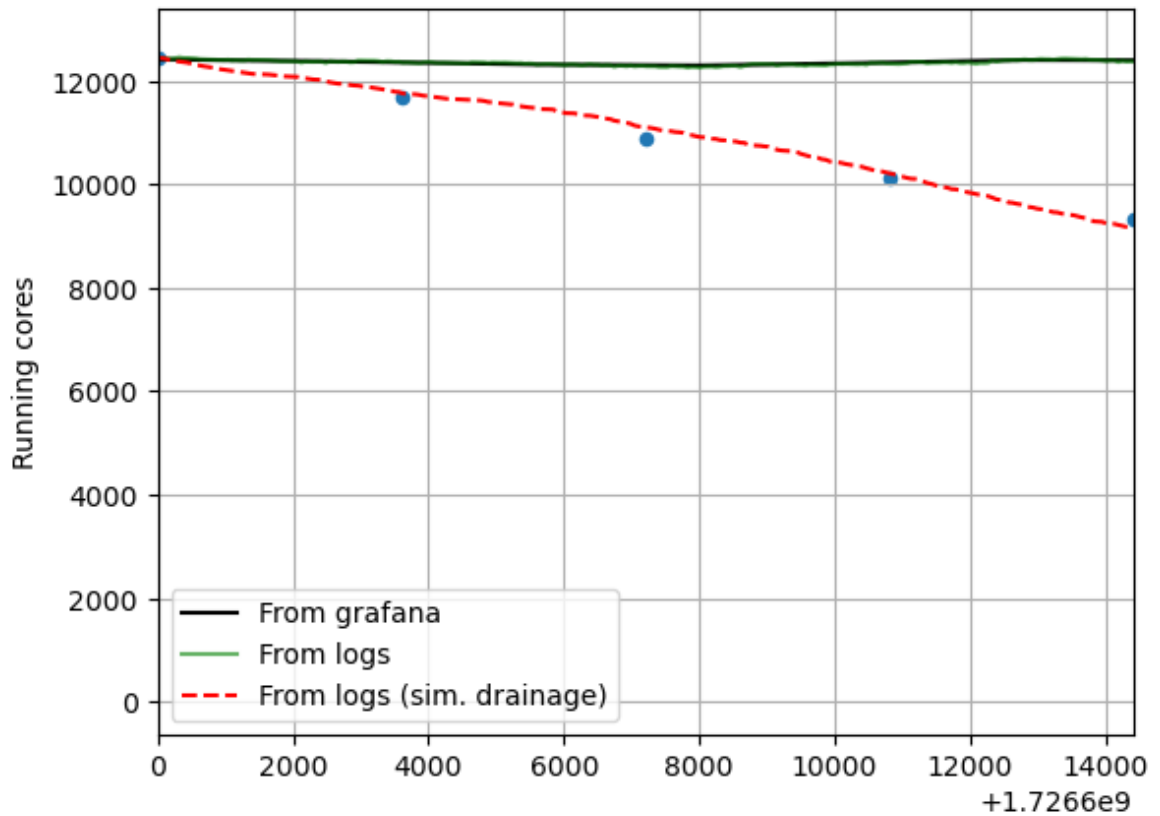
In particular, we observe in this simulation that approx. there is a 25% decrease in farm utilization every 4 hours, or approx. 6% every hour.

```
In [11]: plt.plot(df_graf['Time']/1000, df_graf['All'], label="From grafana", color="black", alpha=0.6)
plt.plot(t_fi, v_fi, label="From logs", color="green", alpha=0.6)
plt.plot(t_fi2, v_fi2, label="From logs (sim. drainage)", color="red", linestyle="dashed", alpha=0.6)

t_red = [time_drop, time_drop+3600, time_drop+2*3600, time_drop+3*3600, time_drop+4*3600]
v_red = [12450, (1-0.25/4)*12450, (1-2*0.25/4)*12450, (1-3*0.25/4)*12450, (1-4*0.25/4)*12450]

plt.scatter(t_red, v_red, s=20)

plt.xlim(time_drop, time_drop+4*3600)
plt.ylabel("Running cores")
plt.legend(loc='best')
plt.grid()
plt.show()
```



Next steps

Conduct extended simulations to observe drainage patterns across varying load conditions

Analyze the impact of job types and VO-specific job durations on natural drainage: how different workloads contribute to natural drainage cycles

Estimate power consumption before and after drainage (ipmitool) - we know which job slots are getting free, and in which machine

Evaluate potential carbon emission reductions in these drainage scenarios

Develop machine learning models for predictive power scaling

Design a feedback loop to HTCondor for real-time power modulation based on green energy cycles