

Introduction

- The FRESCO project, sourcing data from Purdue, the University of Illinois at Urbana-Champaign, and the University of Texas at Austin, focuses on improving computer system dependability through collection and curation of detailed system usage data, workloads, and outages.
- The data includes aspects like job submissions, resource allocation, and durations, essential for understanding computing system failures and utilization.

Questions

- How do jobs utilize cluster resources in university's centrally managed clusters?
- How do users use or do not use the options to share resources on a node?
- How often do resource demands exceed supply, and does this impact job failure rates?
- Can users estimate the time their jobs will need on the cluster?

FRESCO Data Tables

Job Data Table

Column	Type
jid	character varying(32)
submit_time	timestamp with time zone
start_time	timestamp with time zone
end_time	timestamp with time zone
runtime	real
timelimit	real
node_hrs	real
nhosts	integer
incores	integer
ngpus	integer
username	character varying(64)
account	character varying(64)
queue	character varying(64)
state	character varying(64)
jobname	text
exitcode	text
host_list	text[]

Host Data Table

Column	Type
time	timestamp with time zone
host	character varying(64)
jid	character varying(32)
type	character varying(32)
event	character varying(64)
unit	character varying(16)
value	real
diff	real
arc	real

- This table contains accounting information for each job.
- The 'exitcode' column allows us to identify failed jobs.
- This table provides data for each host in the cluster.
- Represented as a timeseries such that each row represents a single value for a given event type.

Data Analysis

Please select a statistic to calculate.

Statistics: None Mean Median Standard Deviation PDF

Please provide the threshold if 'Ratio of Data Outside Threshold' was selected.

Value:

Please select an interval type to use in the statistic calculation. If count is selected, the interval will correspond to a count of rows. If time is selected, the interval will be a time window.

Interval Type:

If time was selected, please select the unit of time.


Interval Unit:

Please provide the interval count.

Value:

- After creating a dataset from the Host Data table, users can perform statistical calculations. The available metrics are:
- cpuuser**: CPU user mode average %
- gpu_usage**: GPU active time average %
- block**: data transfer rate
- memused**: total memory storage
- memused_minus_diskcache**: physical memory usage excluding caches
- nfs**: data transfer rate over NFS mounts

The Anvil System



Data Access

Query the Host Data Table

Select start and end times (Max: 5 days).

Start Time:

End Time:

Times Valid

Select columns:

Columns:

Select Distinct

Choose sort column and direction:

Order By:

Direction:

Set results limit:

Limit Results:

Enter IN clause values:

IN Column:

IN values:

Add data filters:

Column:

Operator:

Value:

Add Condition

Active filters:

Conditions:

Execute Query

Current SQL query:

```
SELECT * FROM host_data WHERE time BETWEEN %s AND %s
Parameters: [datetime.datetime(2022, 10, 30, 19, 2, 26), datetime.datetime(2022, 12, 30, 20, 2, 26)]
```

Query the Job Data Table

Select start and end times (Max: 180 days).

Start Time:

End Time:

Times Valid

Select columns:

Columns:

Select Distinct

Choose sort column and direction:

Order By:

Direction:

Set results limit:

Limit Results:

Enter IN clause values:

IN Column:

IN values:

Add data filters:

Column:

Operator:

Value:

Add Condition

Active filters:

Conditions:

Execute Query

Current SQL query:

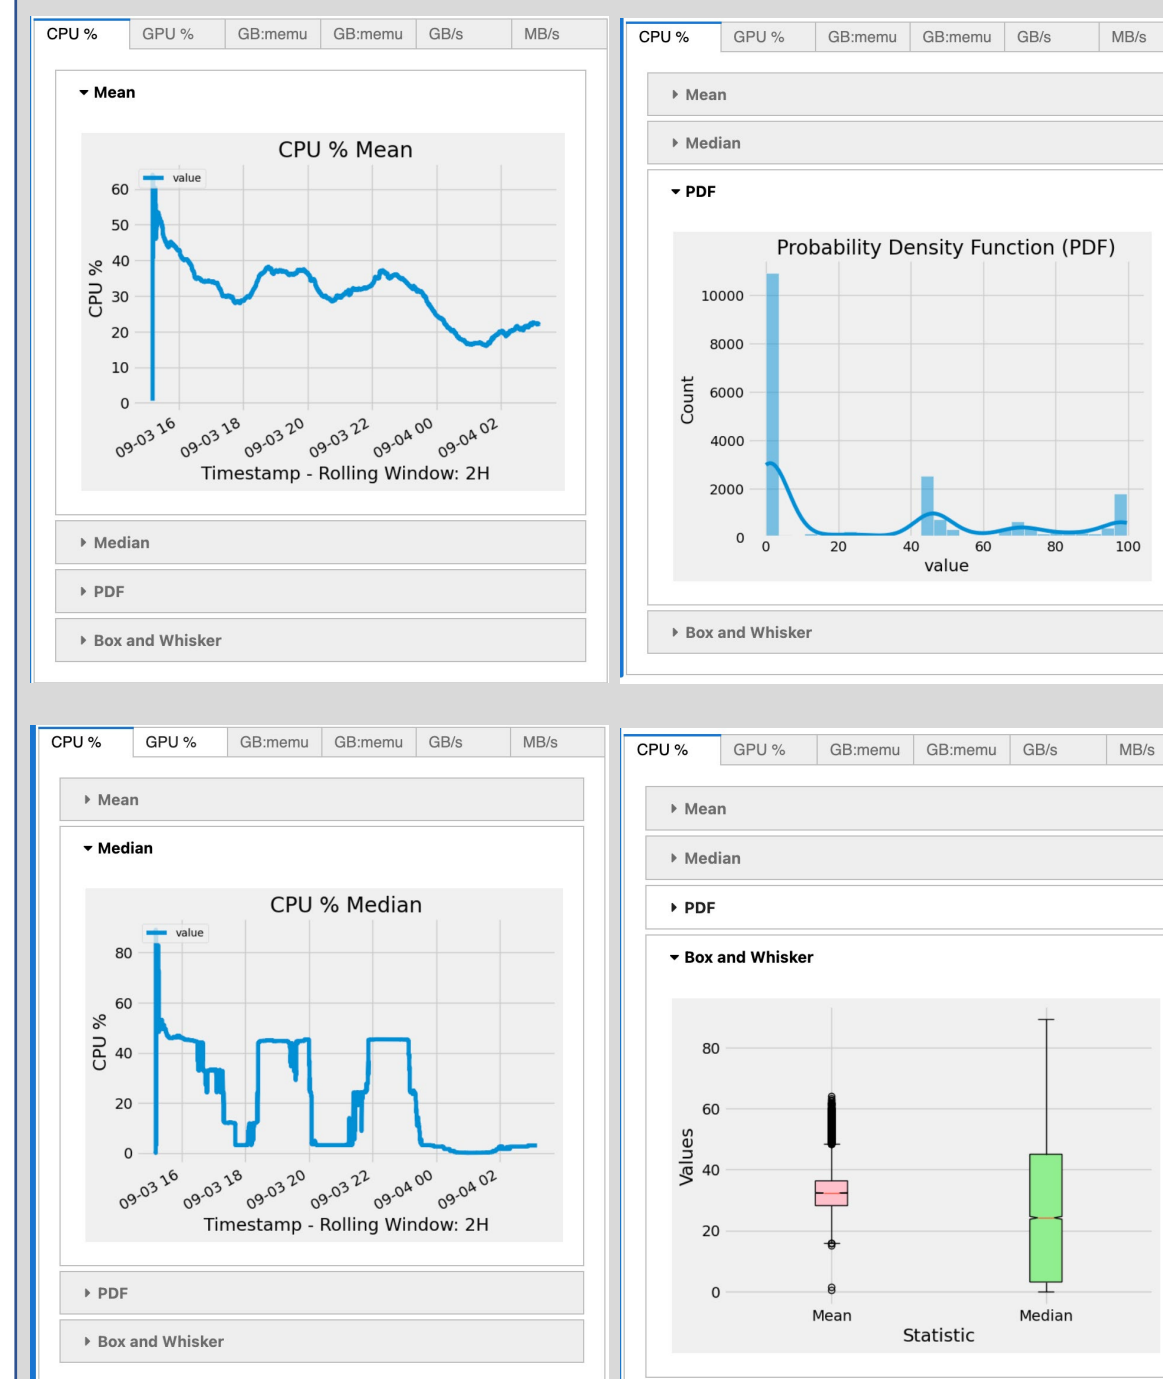
```
SELECT * FROM job_data WHERE start_time BE
TWEEN %s AND %s
Parameters: [datetime.datetime(2022, 10, 30, 19, 2, 26), datetime.datetime(2022, 12, 30, 20, 2, 26)]
```

- Users can interactively request data using a SQL query builder
- Queried data can be exported as CSV or Excel spreadsheet for further external analysis

Data Overview


- Data comes from Anvil, an HPC cluster at Purdue – CPU/GPU/Large Memory jobs
- Data currently ranges from July 2022 to June 2023
- 1,469,223 total jobs; 302,096 failed or timed out jobs
- Tracks job lifecycle events including submission, start, and end times along with exit codes
- Provides detailed node-level resource usage metrics and host event data

Data Visualizations



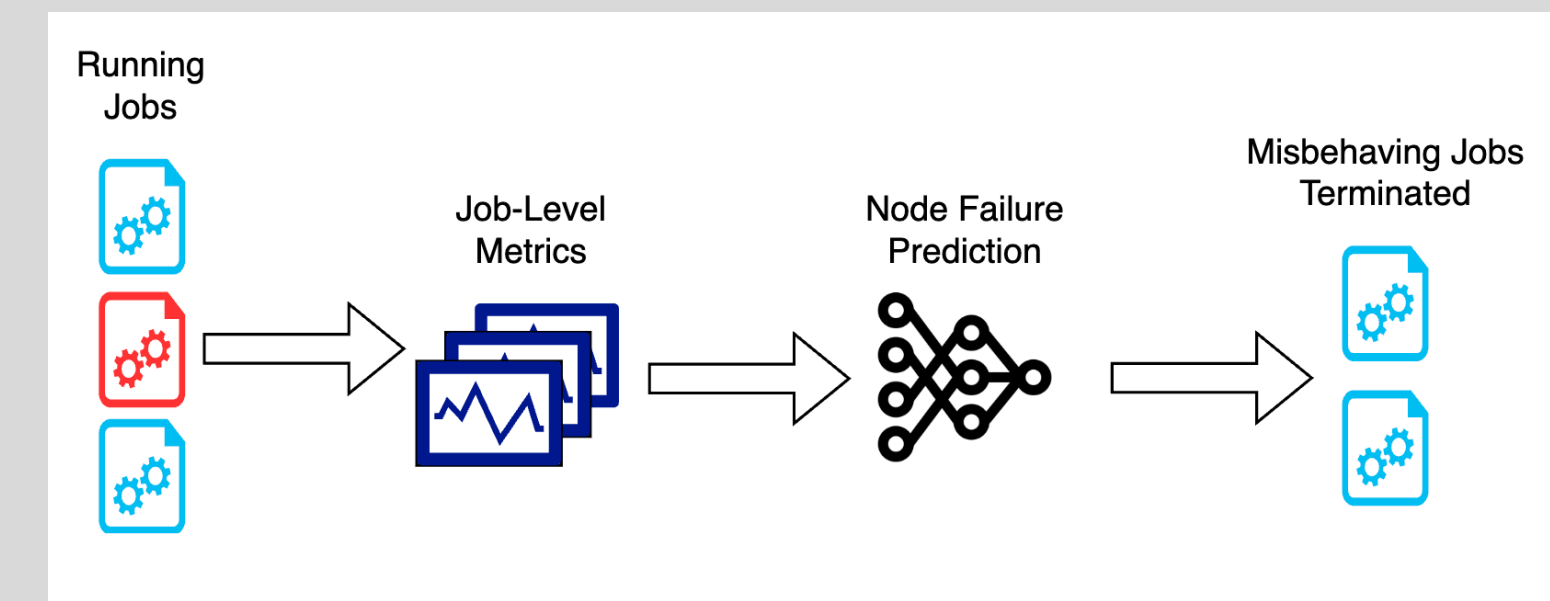
- Plots can be generated interactively as the user queries different subsets of the dataset

Sample Findings



- This plot shows the mean memory usage across jobs on the Z axis, with the measurement index (relative to the measurements for each job) and the walltime of jobs on the X and Y axes respectively
- Blue points are successful jobs, red points are failed jobs
- Anomaly detection methods may be able to learn similar patterns across metrics to predict failed jobs

Future Plans



- Develop a real-time monitoring solution that performs online inference with the node failure prediction model
- Explore using the real-time predictions of the models to terminate jobs that are likely to cause a node failure, preventing the loss of other jobs running on the same node
- Develop models to predict job failures and walltimes, providing quality of service improvements for cluster users

Acknowledgements

- FRESCO is supported by the National Science Foundation, CISE Community Research Infrastructure (CCRI) program, "Open Computer System Usage Repository and Analytics Engine". Project numbers CNS-2016704 and CNS-2016608.
- Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.